

Analysis of Origin Destination Flows (Raw Data)

Anukool Lakhina, Konstantina Papagiannaki, Mark Crovella,
Christophe Diot, Eric D. Kolaczyk, and Nina Taft *

November 10, 2003

BUCS-TR-2003-022

Abstract

In a recent paper, *Structural Analysis of Network Traffic Flows*, we analyzed the set of Origin Destination traffic flows from the Sprint-Europe and Abilene backbone networks. This report presents the complete set of results from analyzing data from both networks. The results in this report are specific to the Sprint-1 and Abilene datasets studied in the above paper. The following results are presented here:

1	Rows of Principal Matrix (V)	2
1.1	Sprint-1 Dataset	2
1.2	Abilene Dataset	9
2	Set of Eigenflows	14
2.1	Sprint-1 Dataset	14
2.2	Abilene Dataset	21
3	Classifying Eigenflows	26
3.1	Sprint-1 Dataset	26
3.2	Abilene Dataset	44

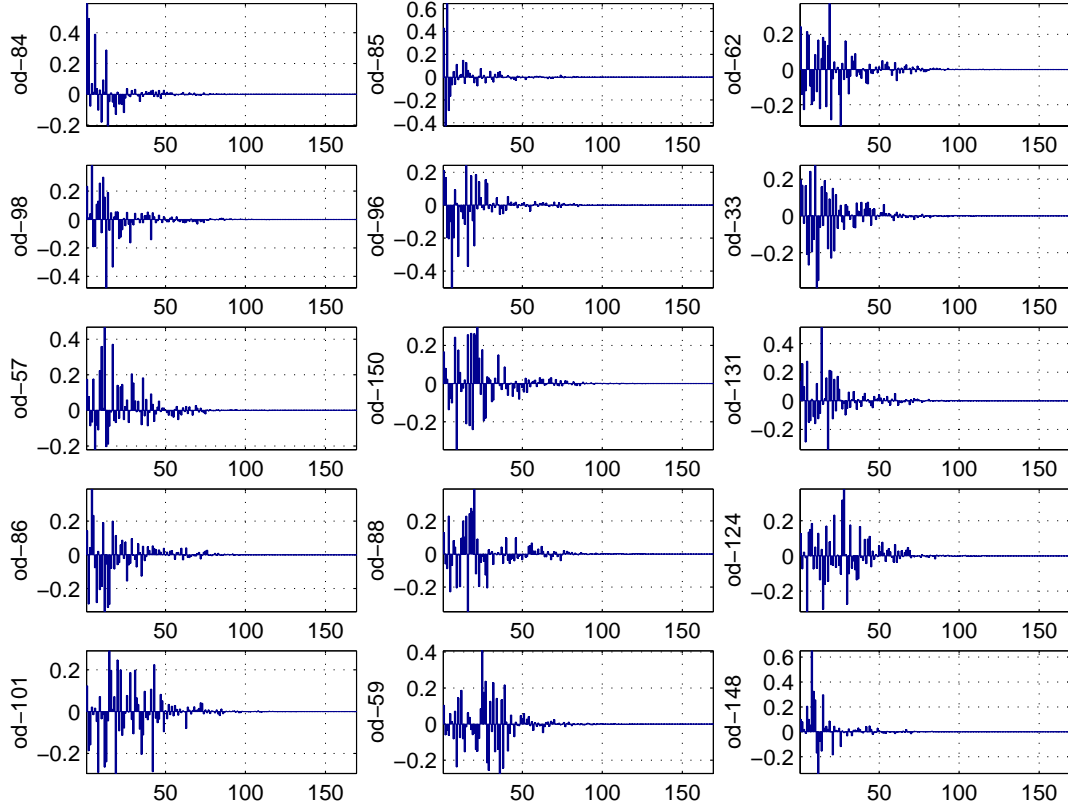
*A. Lakhina and M. Crovella are with the Department of Computer Science, Boston University; email: {anukool,crovella}@cs.bu.edu. K. Papagiannaki is with Sprint Advanced Technology Labs; email: dina@sprintlabs.com. C. Diot is with Intel Research, Cambridge; email: christophe.diot@intel.com. E. D. Kolaczyk is with the Department of Mathematics and Statistics, Boston University; email: kolaczyk@math.bu.edu. N. Taft is with Intel Research, Berkeley; email: nina.taft@intel.com. This work was performed while M. Crovella was at Laboratoire d'Informatique de Paris 6 (LIP6), with support from Centre National de la Recherche Scientifique (CNRS) France. Part of this work was done when N. Taft was at Sprint Advanced Technology Labs and A. Lakhina was at Sprint Advanced Technology Labs and Intel Research, Cambridge. This work was in part supported by a grant from Sprint Labs, ONR award N000140310043 and NSF grants ANI-9986397 and CCR-0325701.

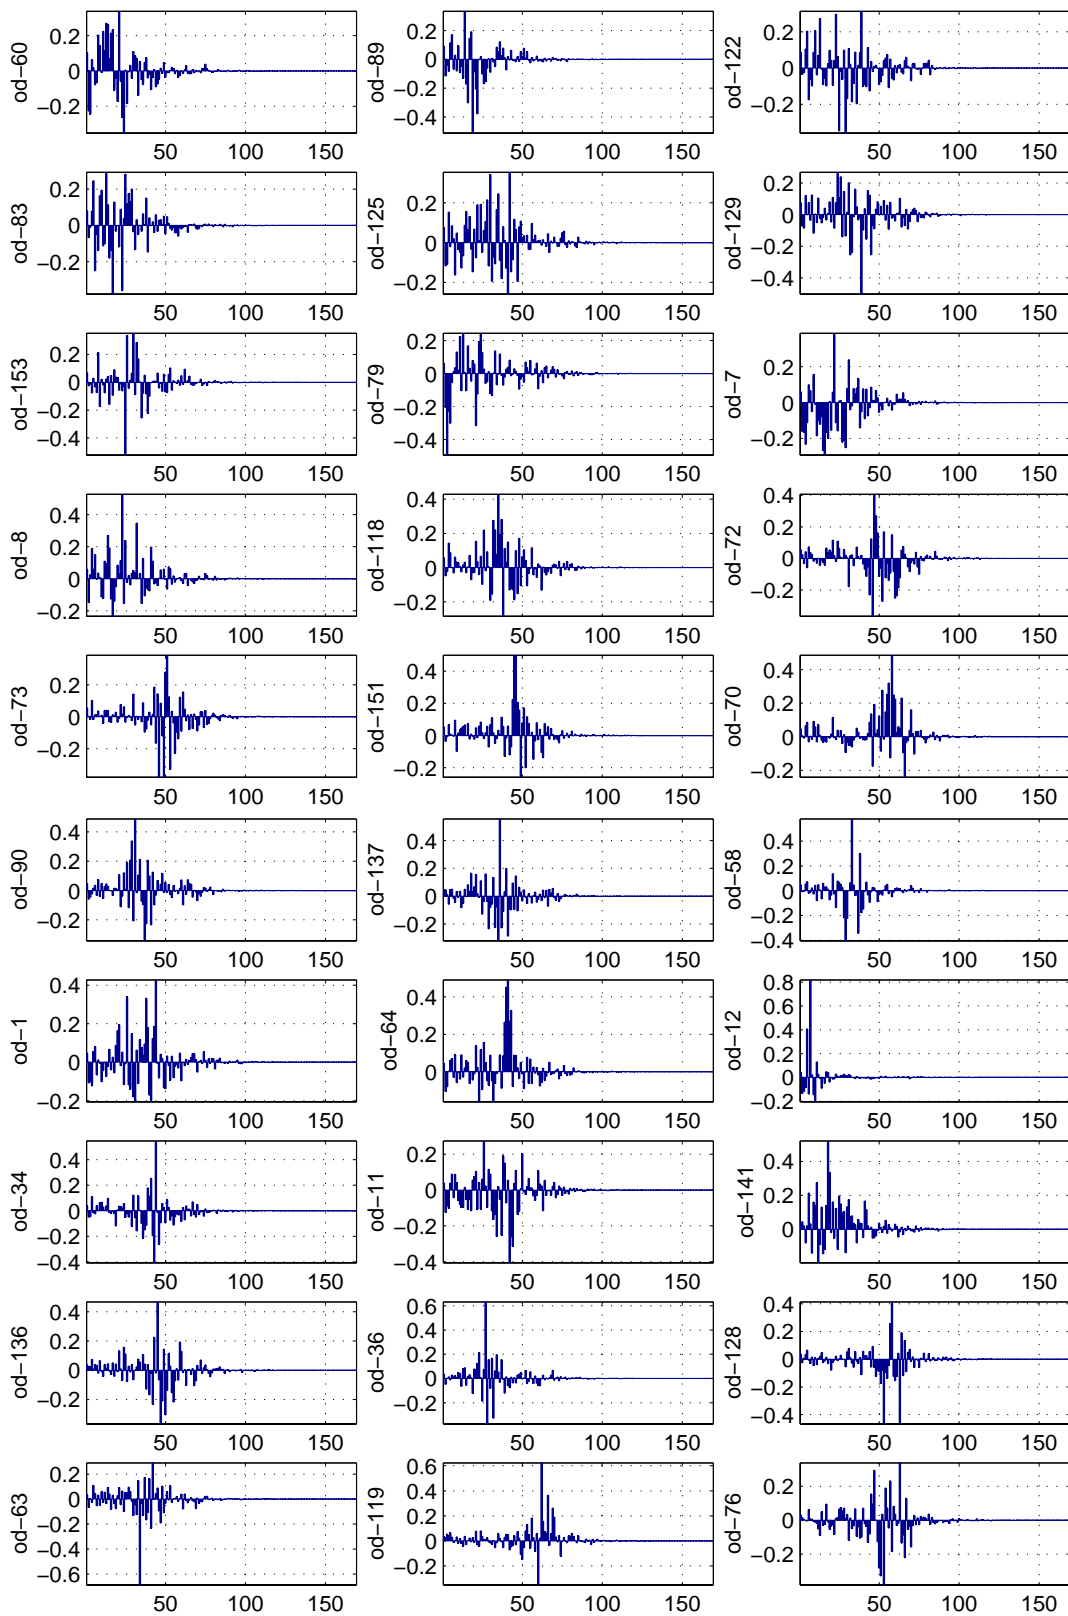
Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 10 NOV 2003		2. REPORT TYPE		3. DATES COVERED 00-00-2003 to 00-00-2003	
4. TITLE AND SUBTITLE Analysis of Origin Destination Flows (Raw Data)				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Boston University, Department of Computer Science, 111 Cummington Street, Boston, MA, 02215				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES The original document contains color images.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 56	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

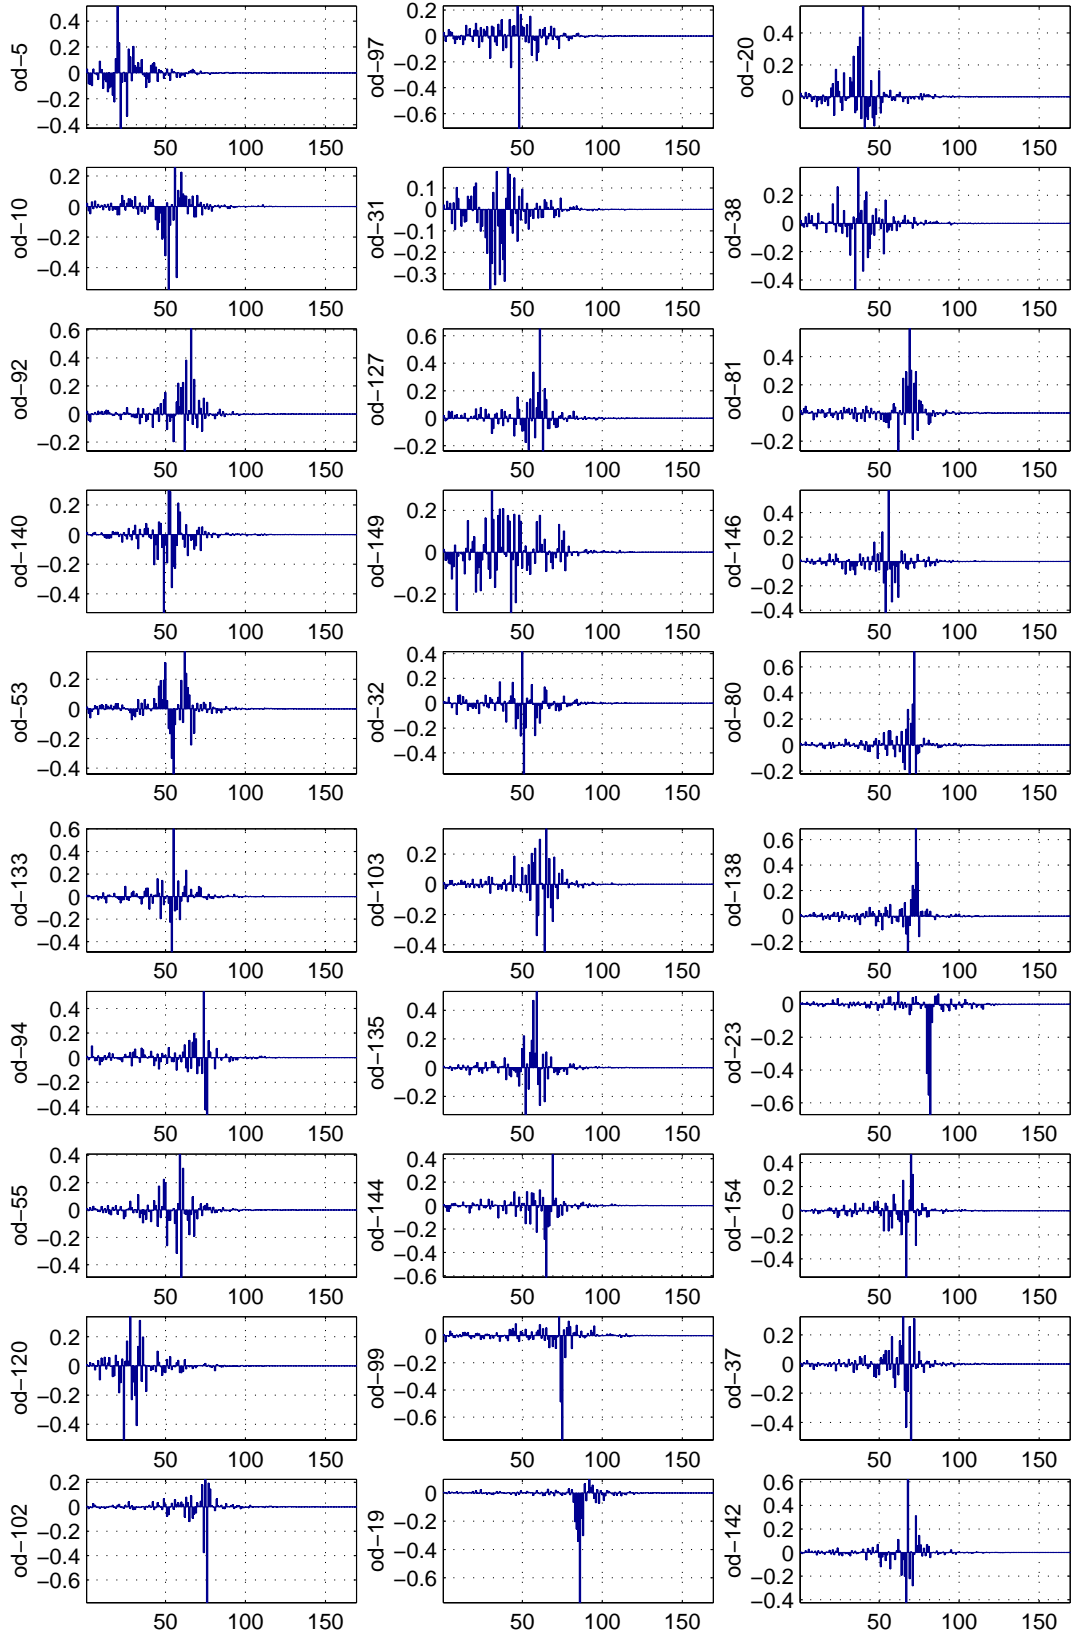
1 Rows of Principal Matrix (V)

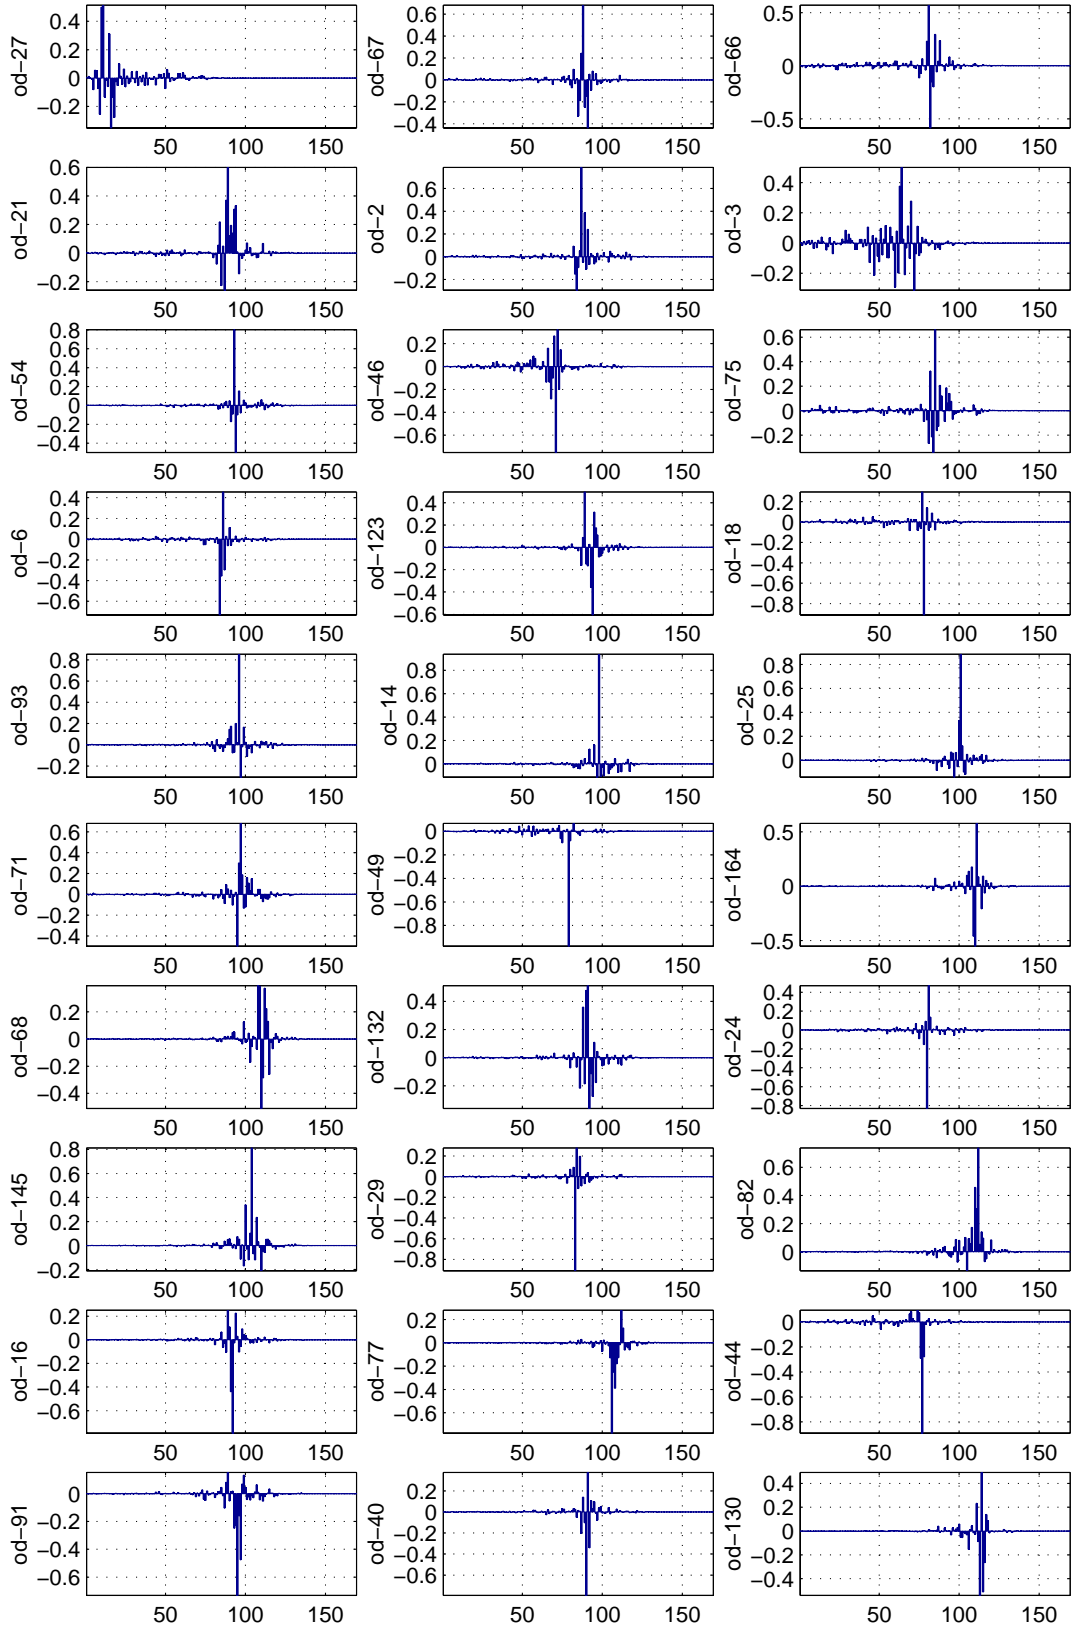
In this section, we present all the rows of principal matrix V for the Sprint-1 and Abilene datasets. Please note that the follow results are sorted from large to small OD flow (in terms of mean rate). And the ylabel corresponding to each plot is the OD flow ID.

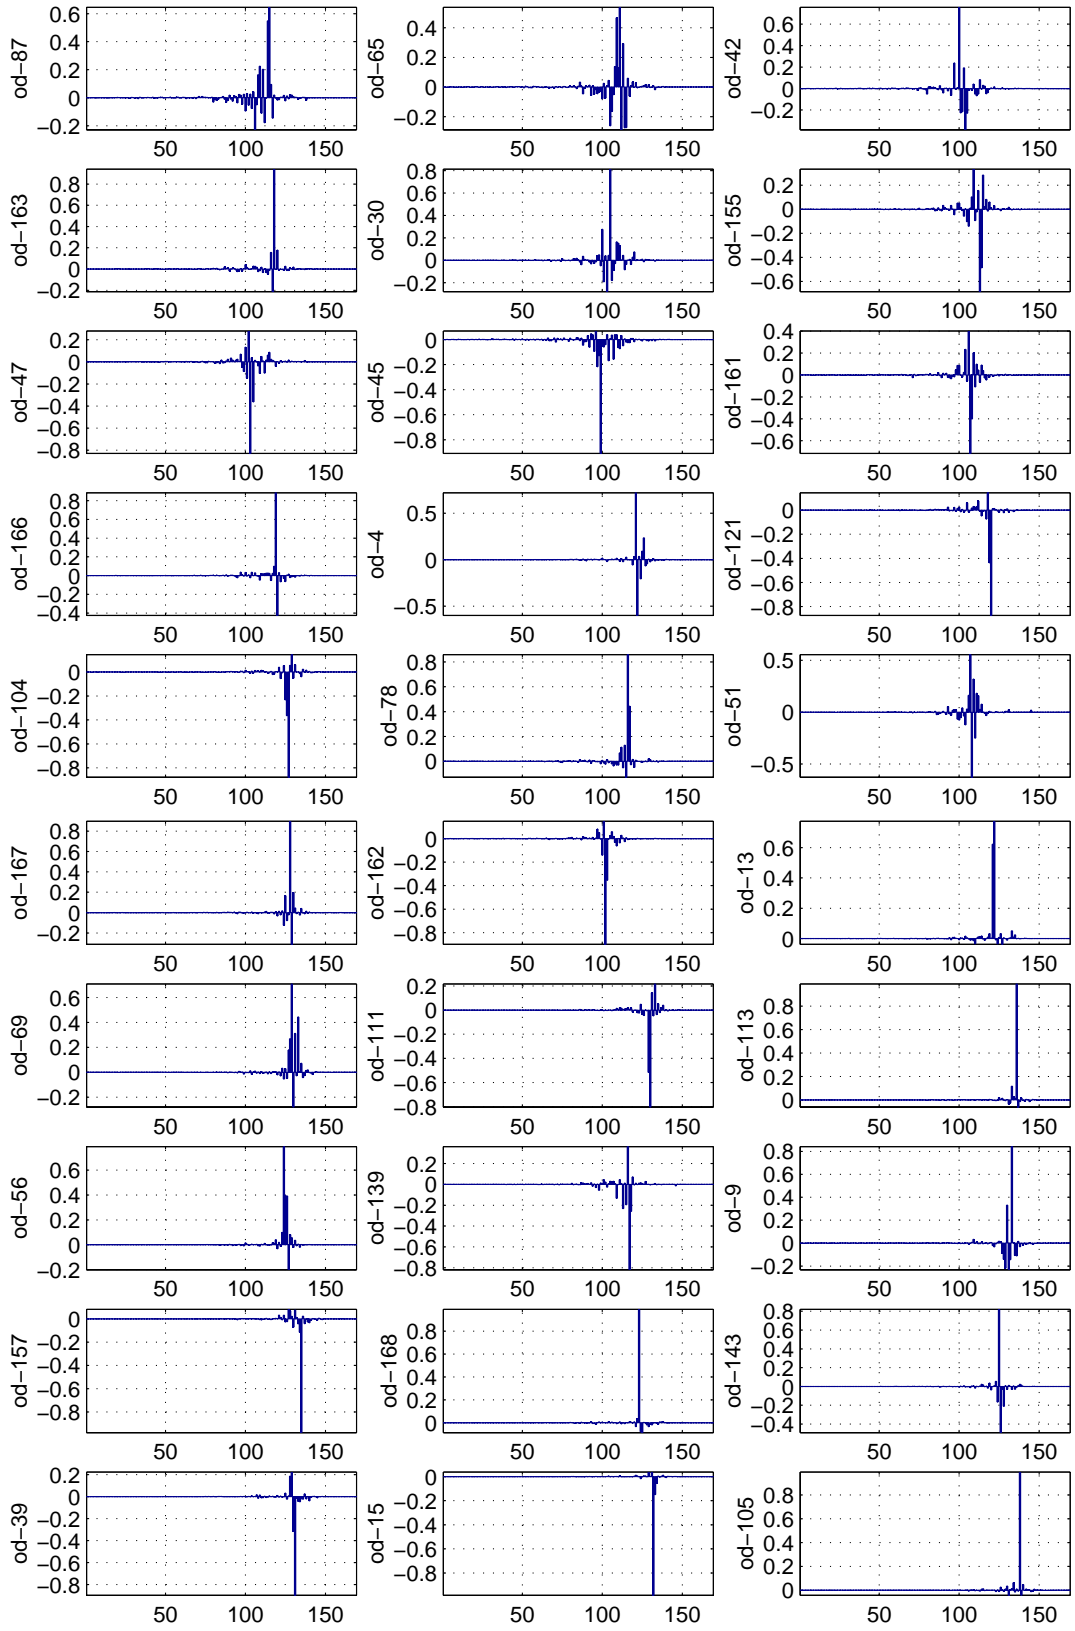
1.1 Sprint-1

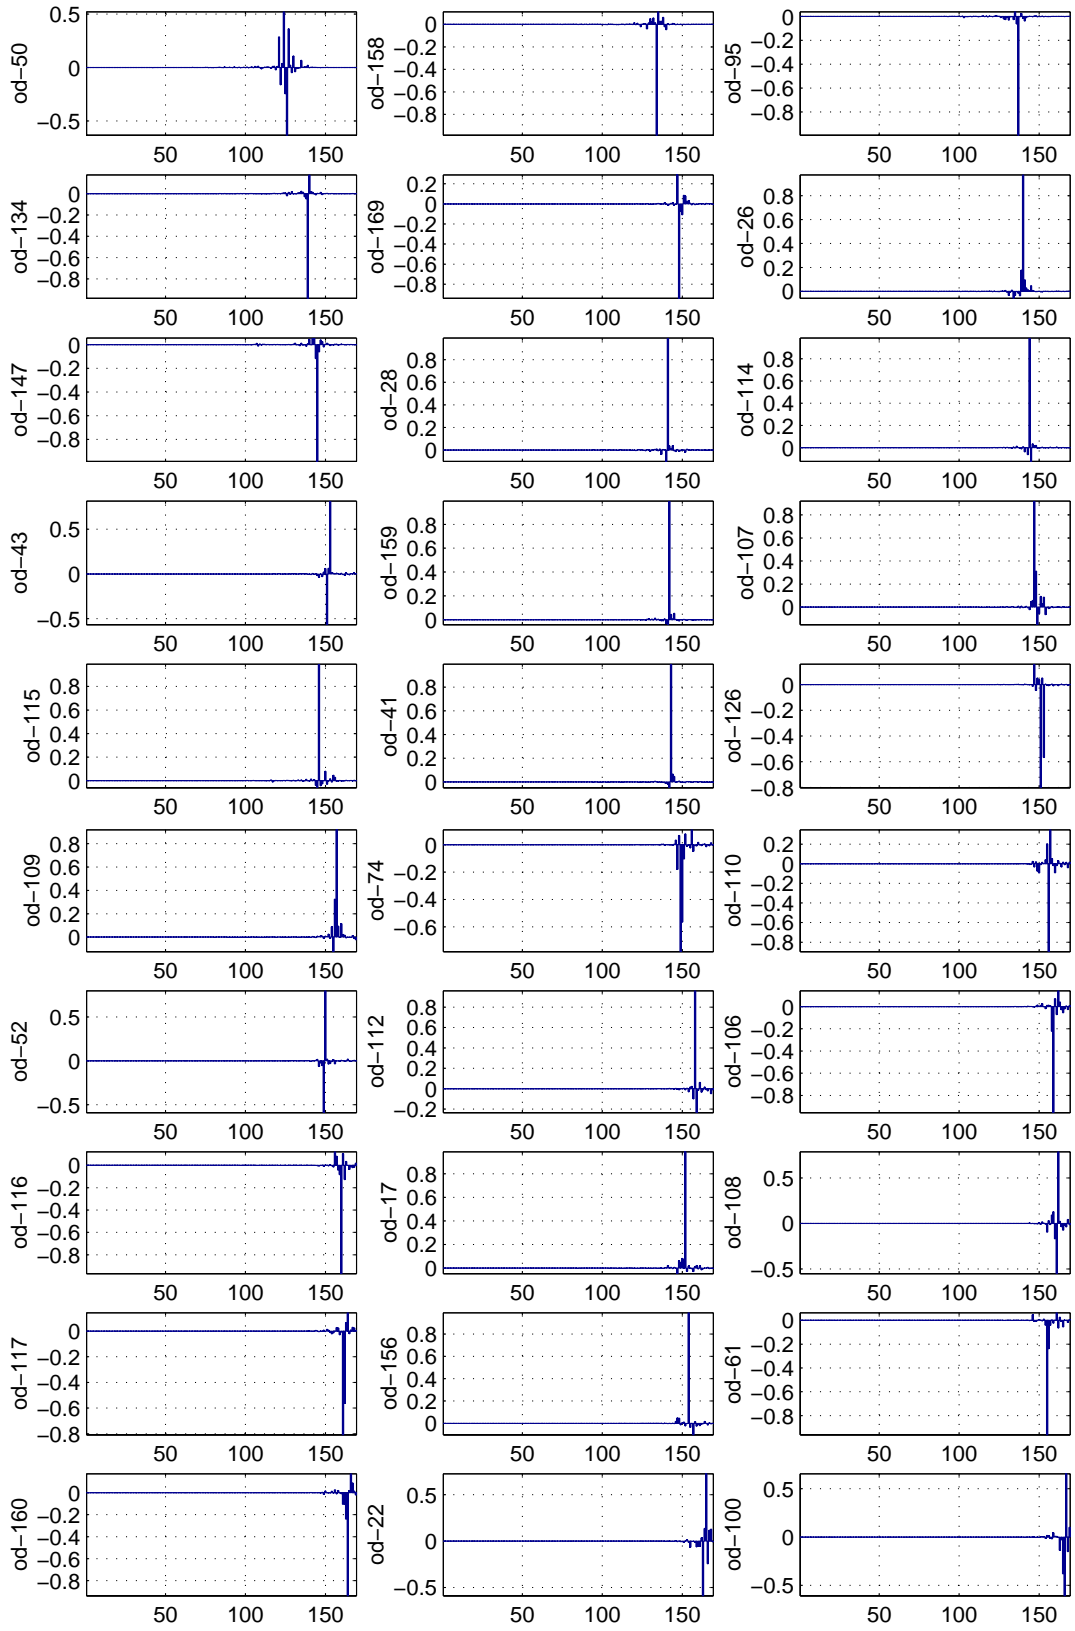










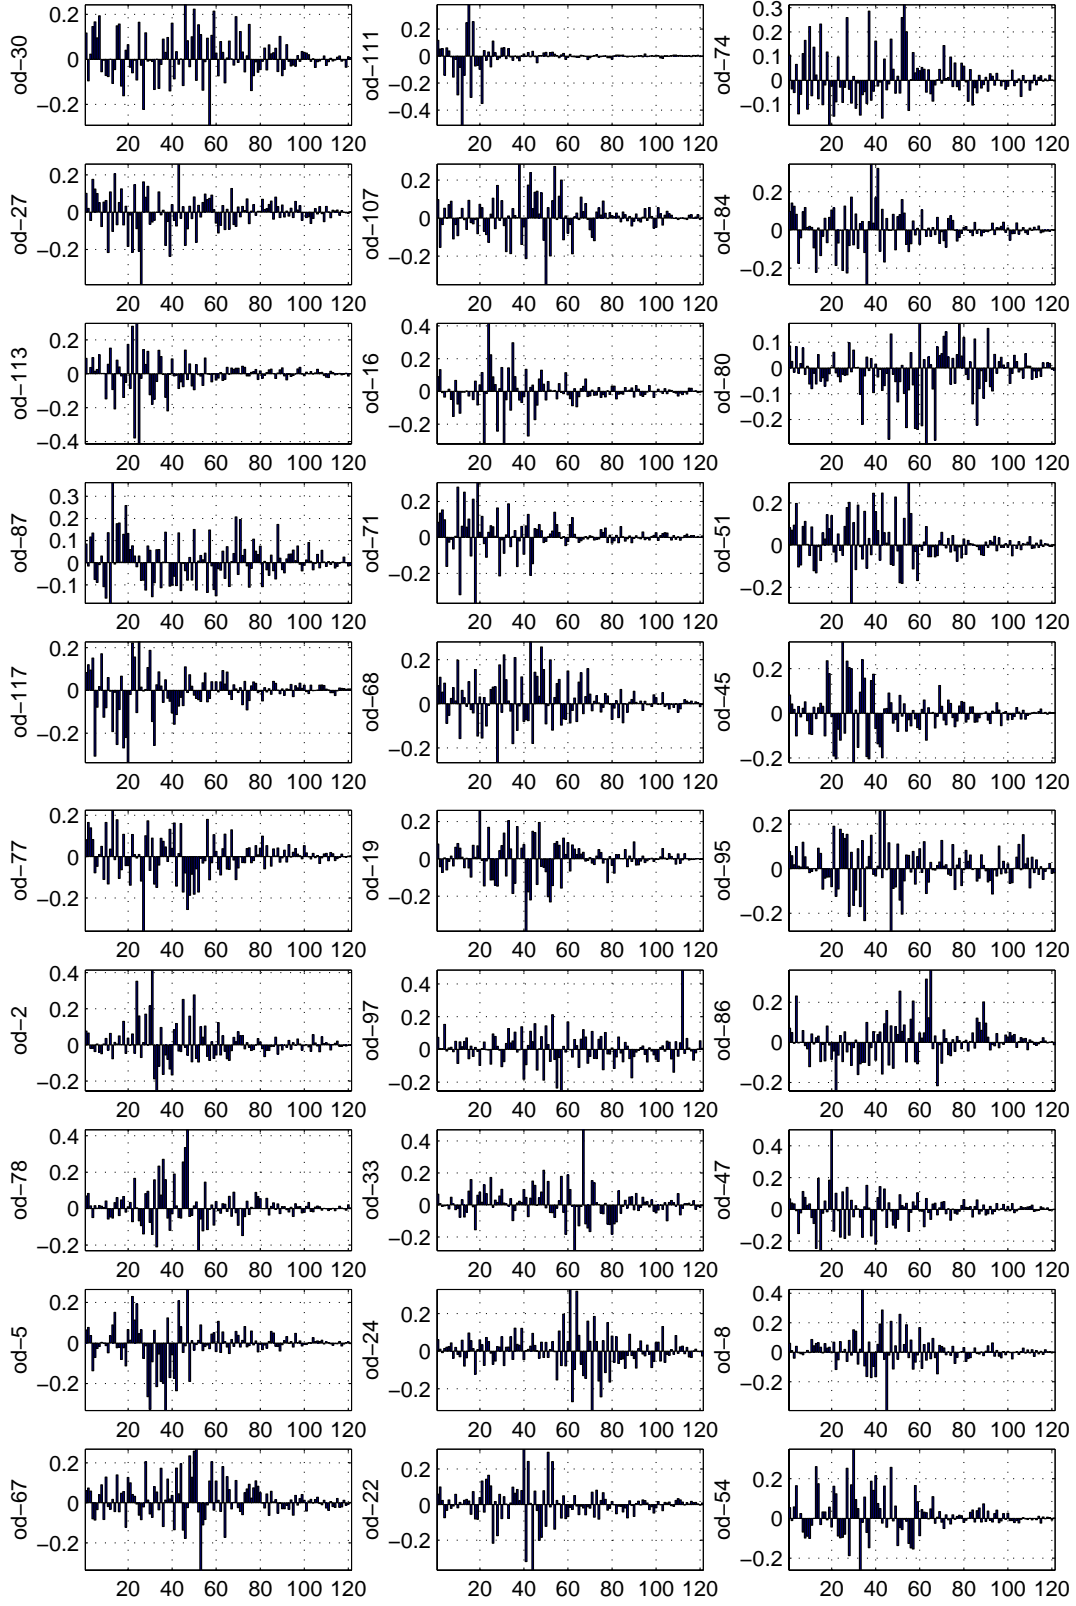


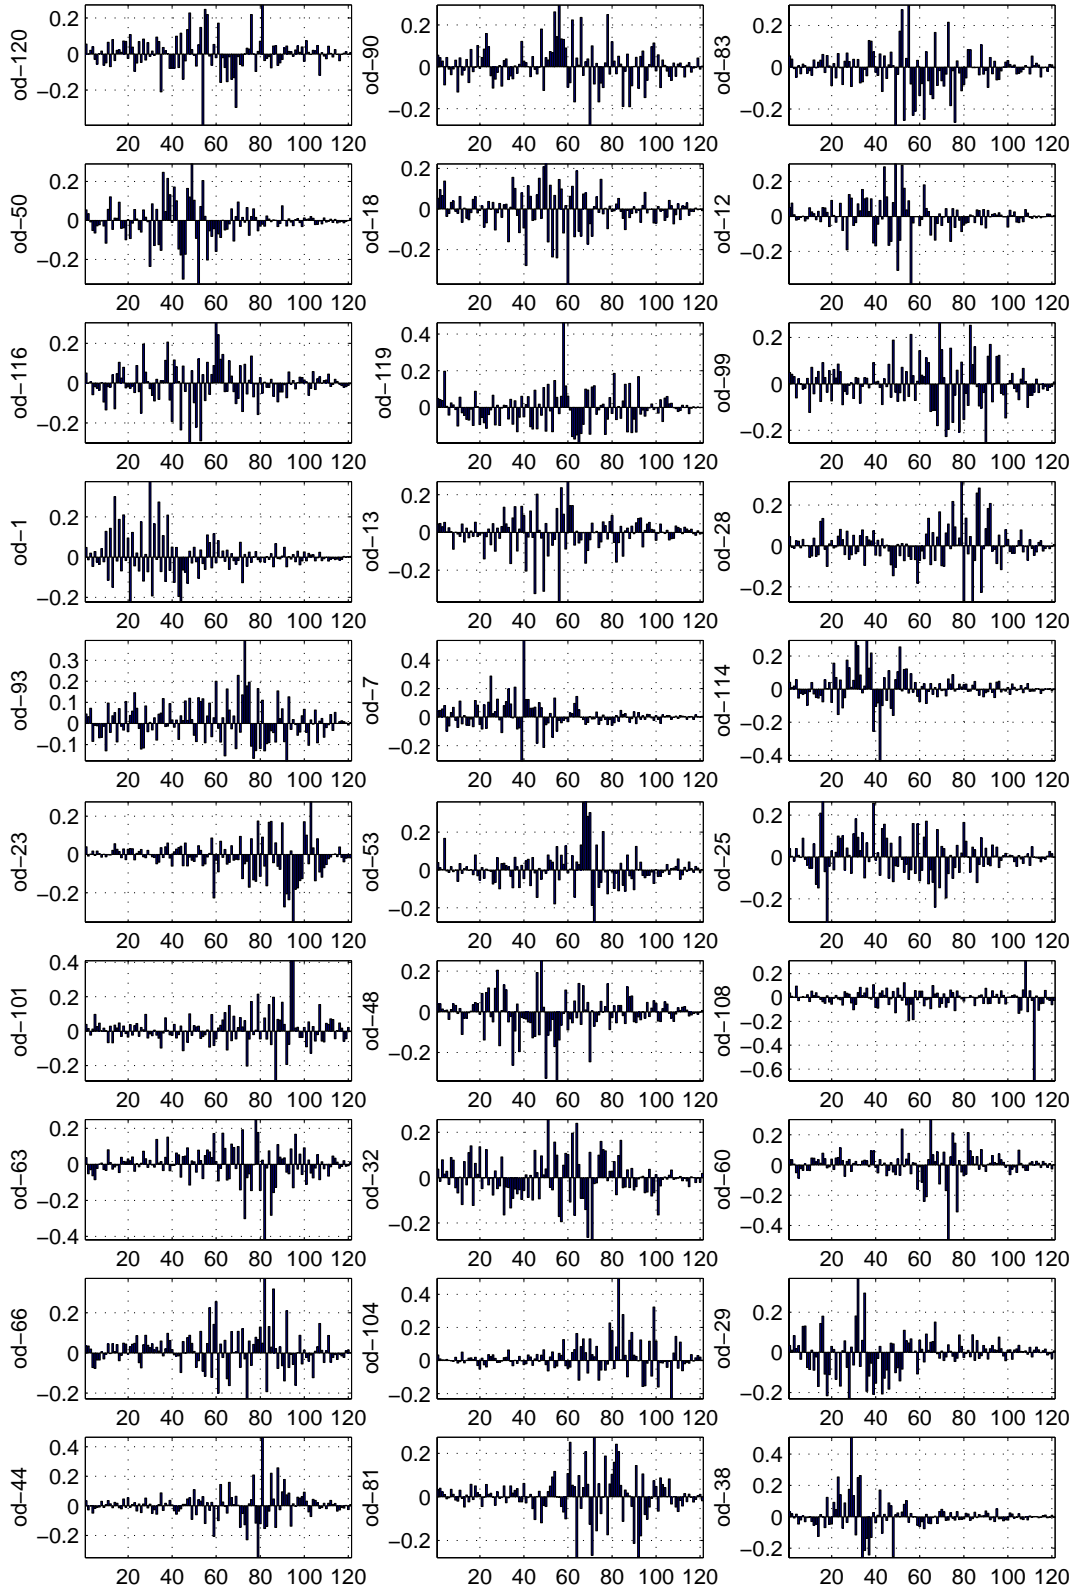


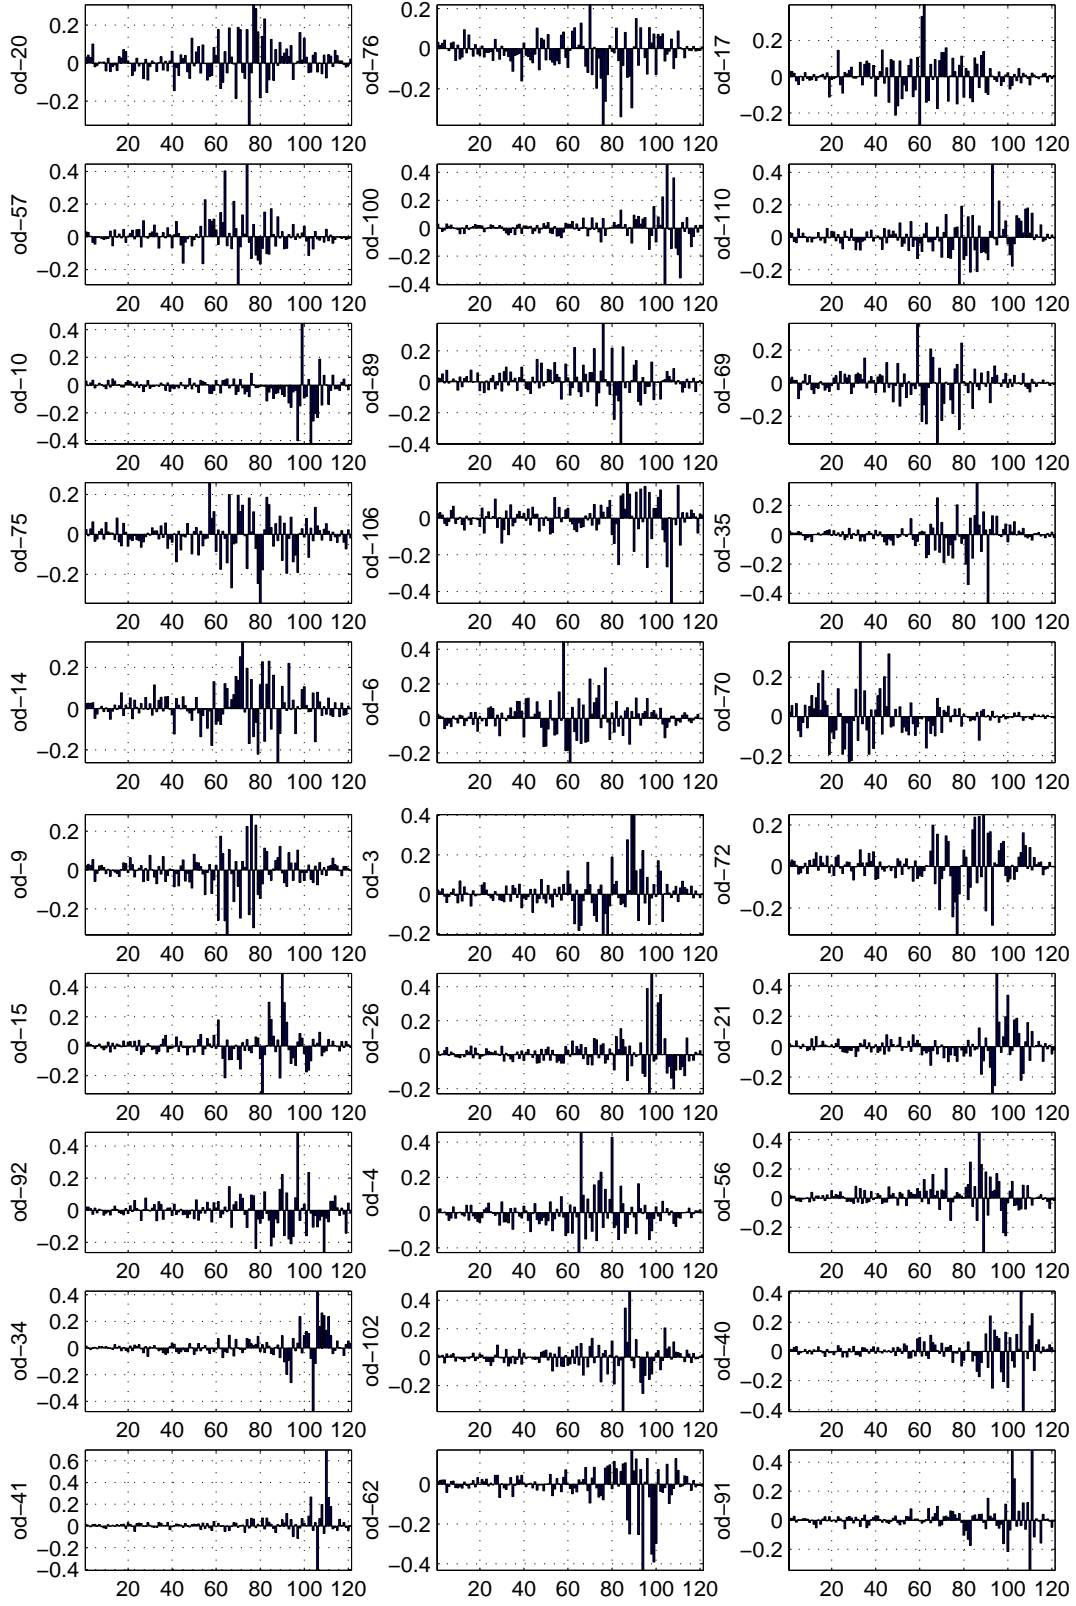


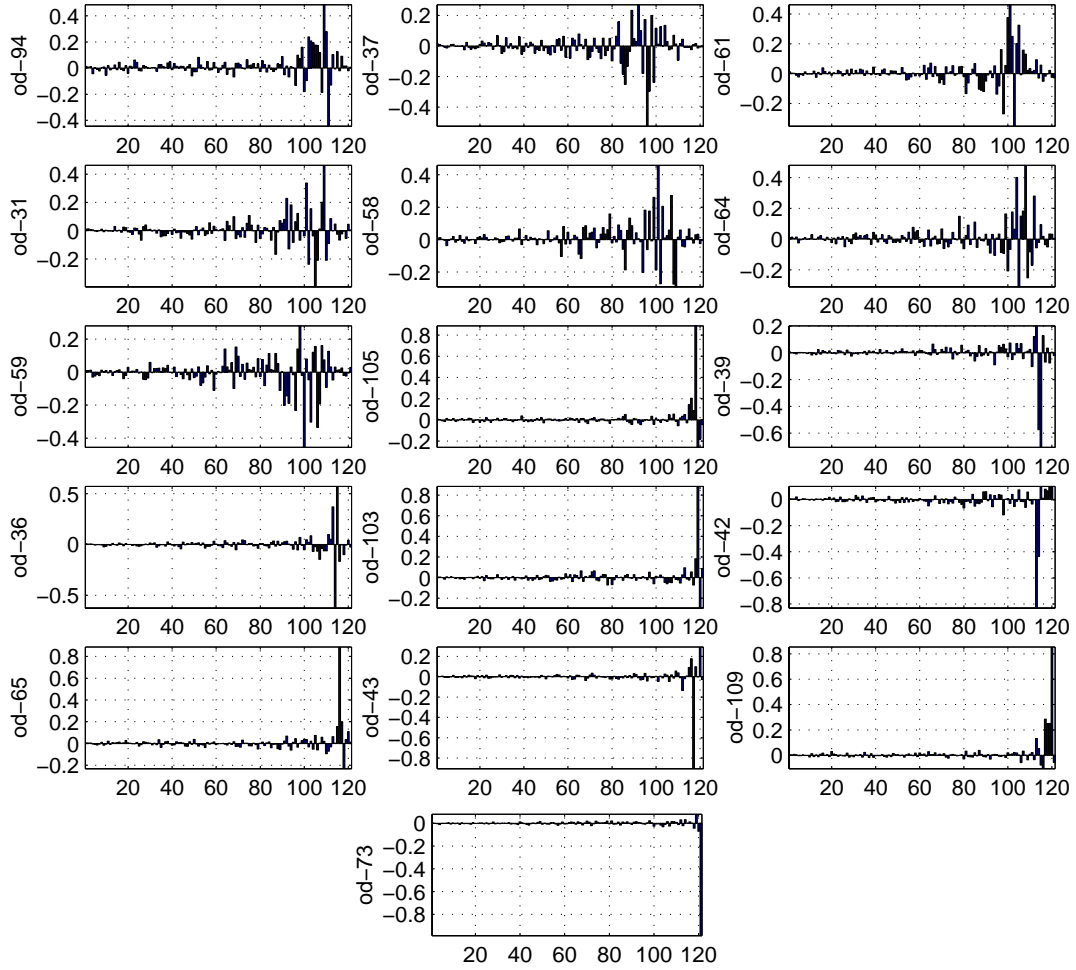
1.2 Abilene







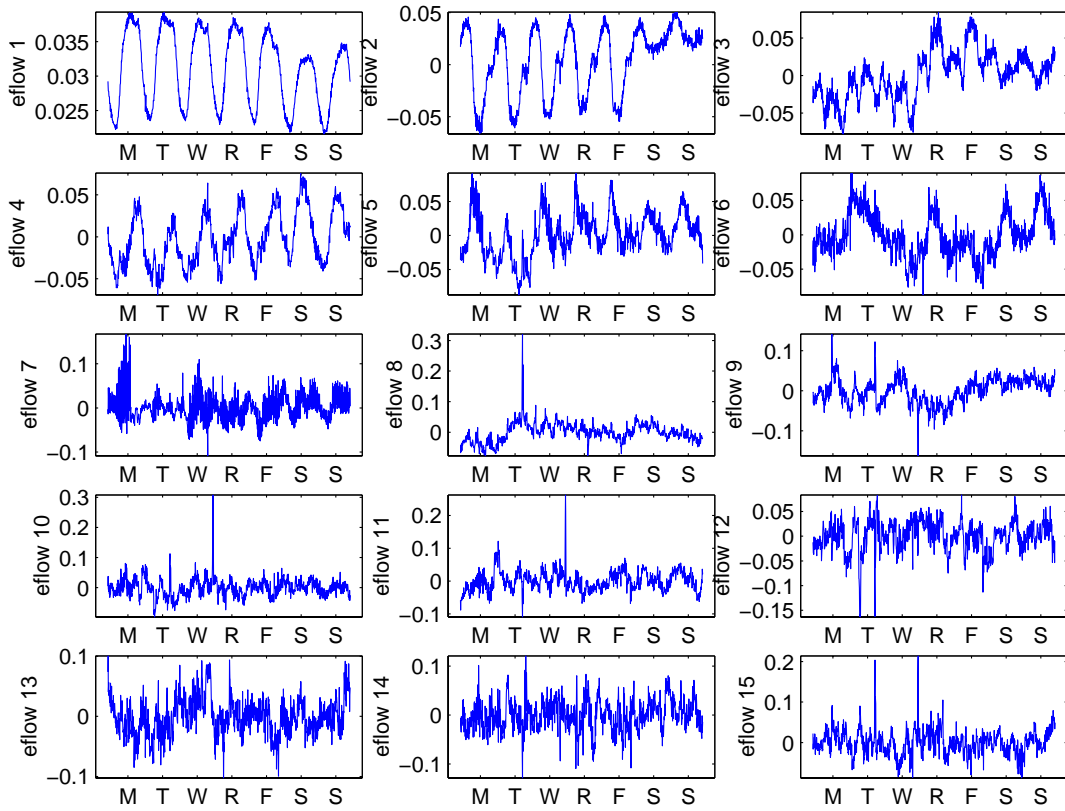


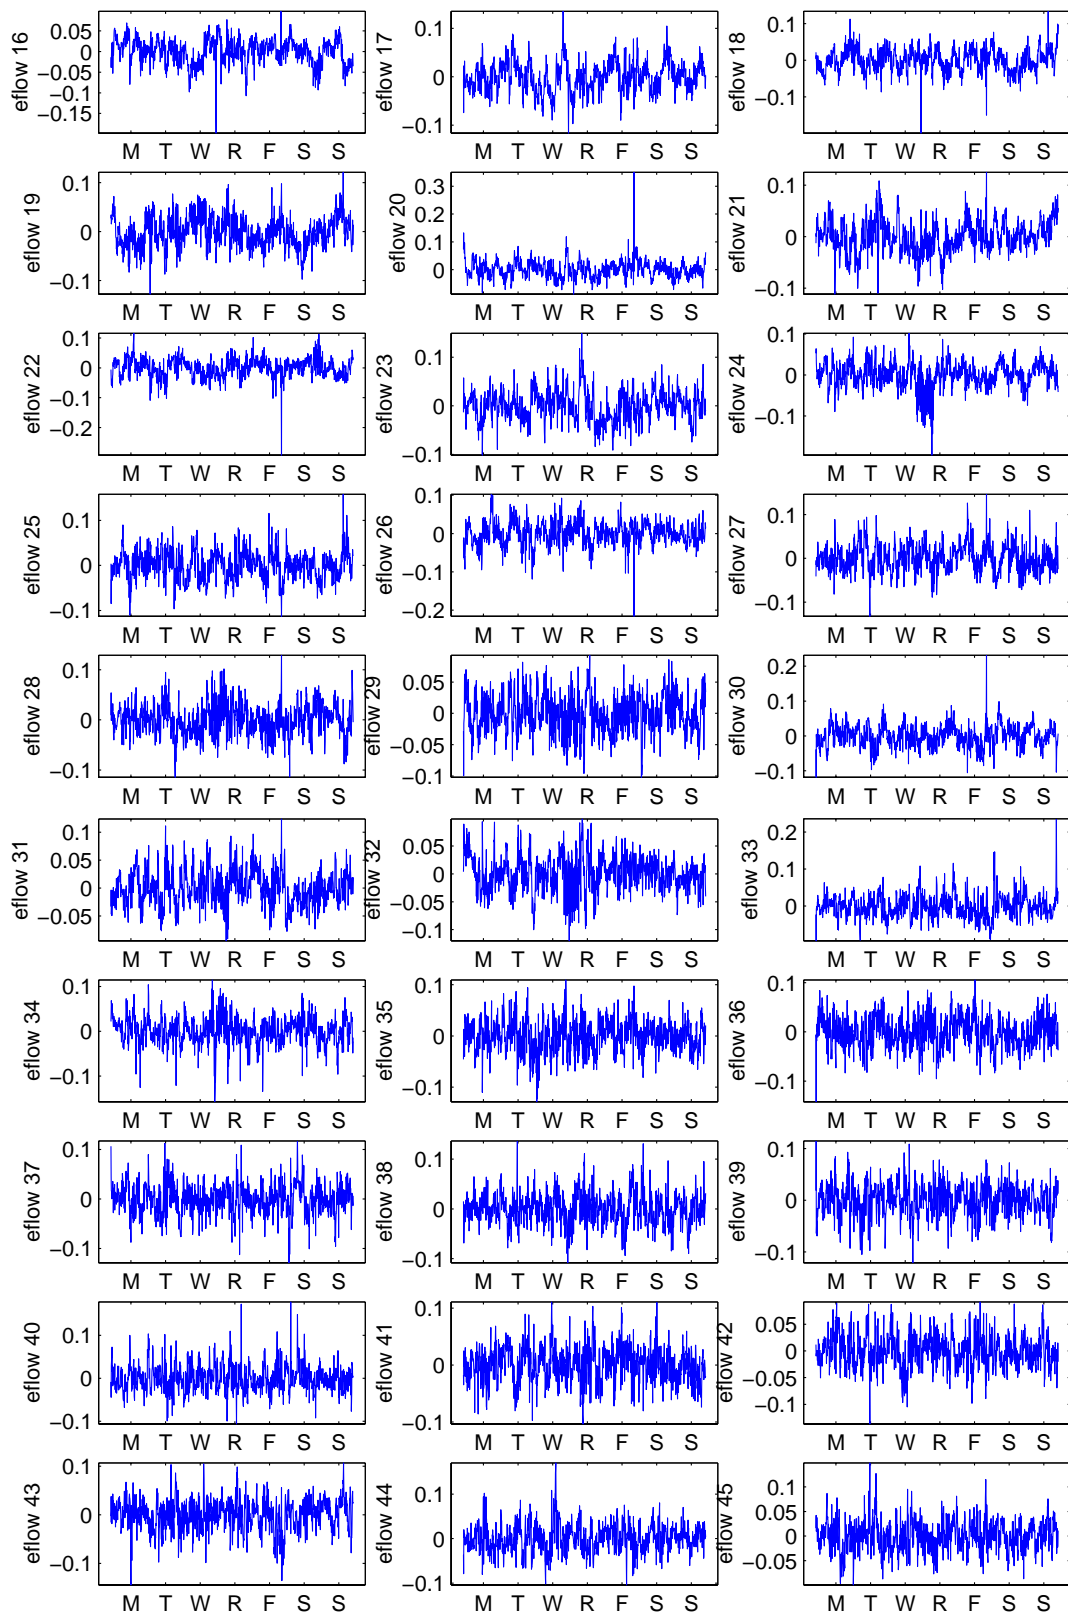


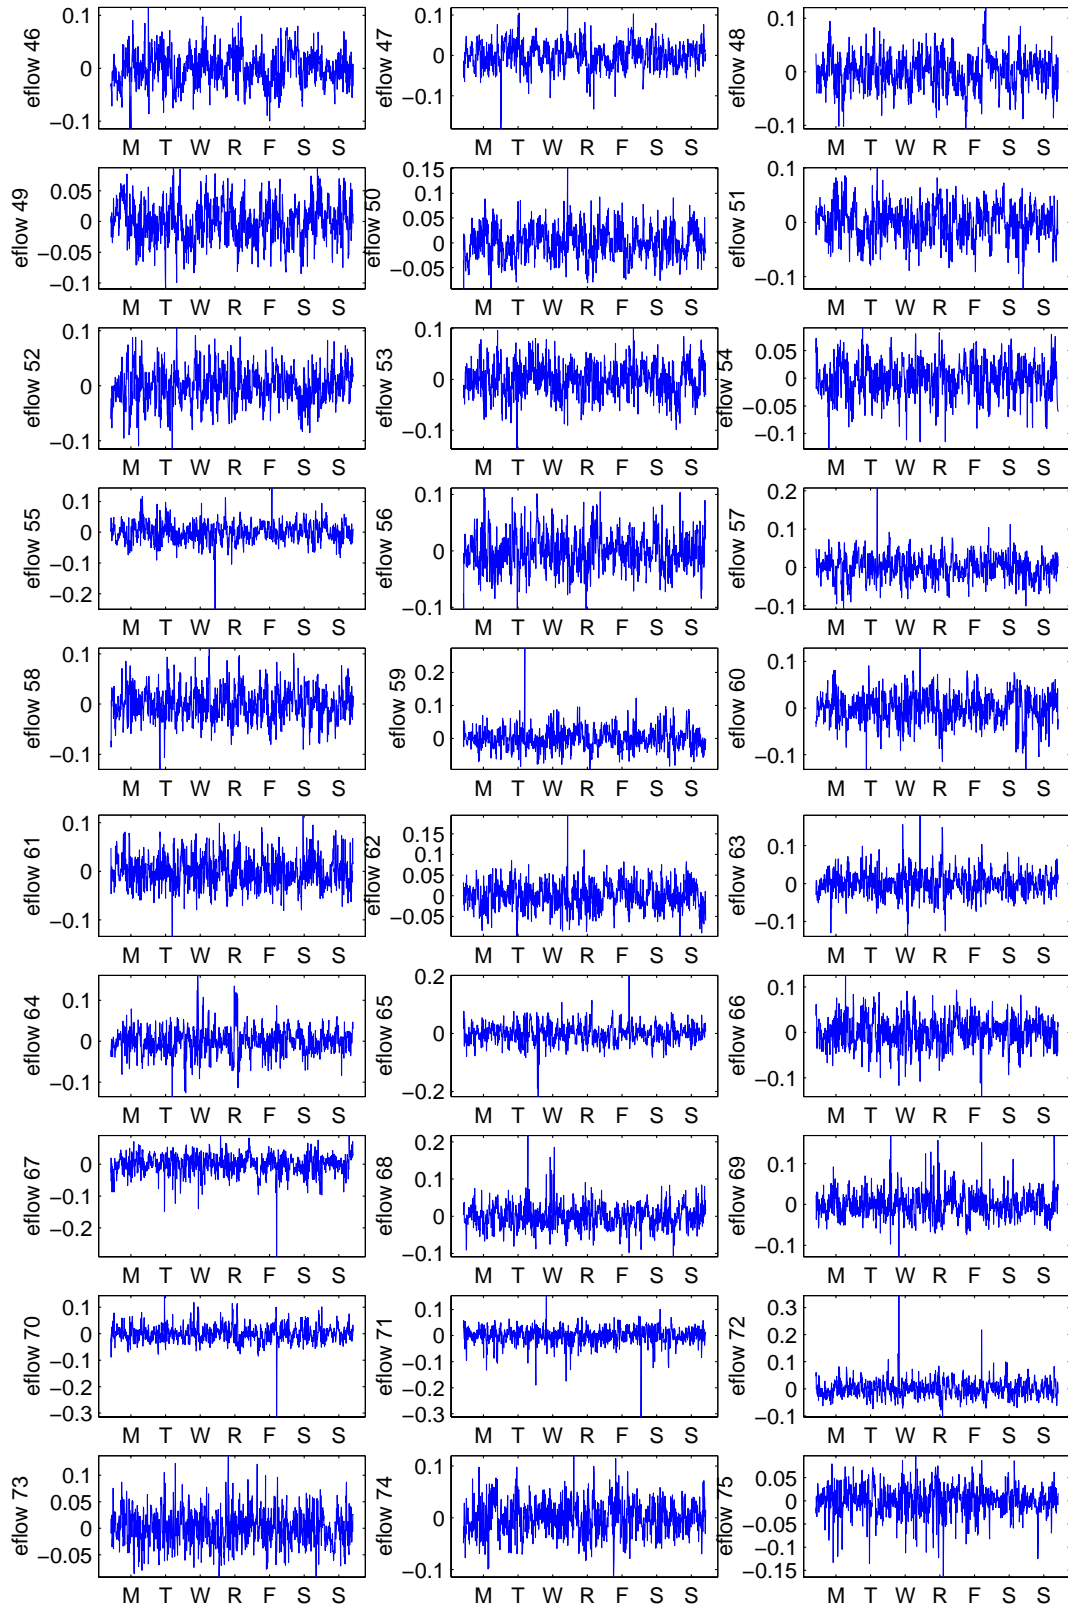
2 Eigenflows

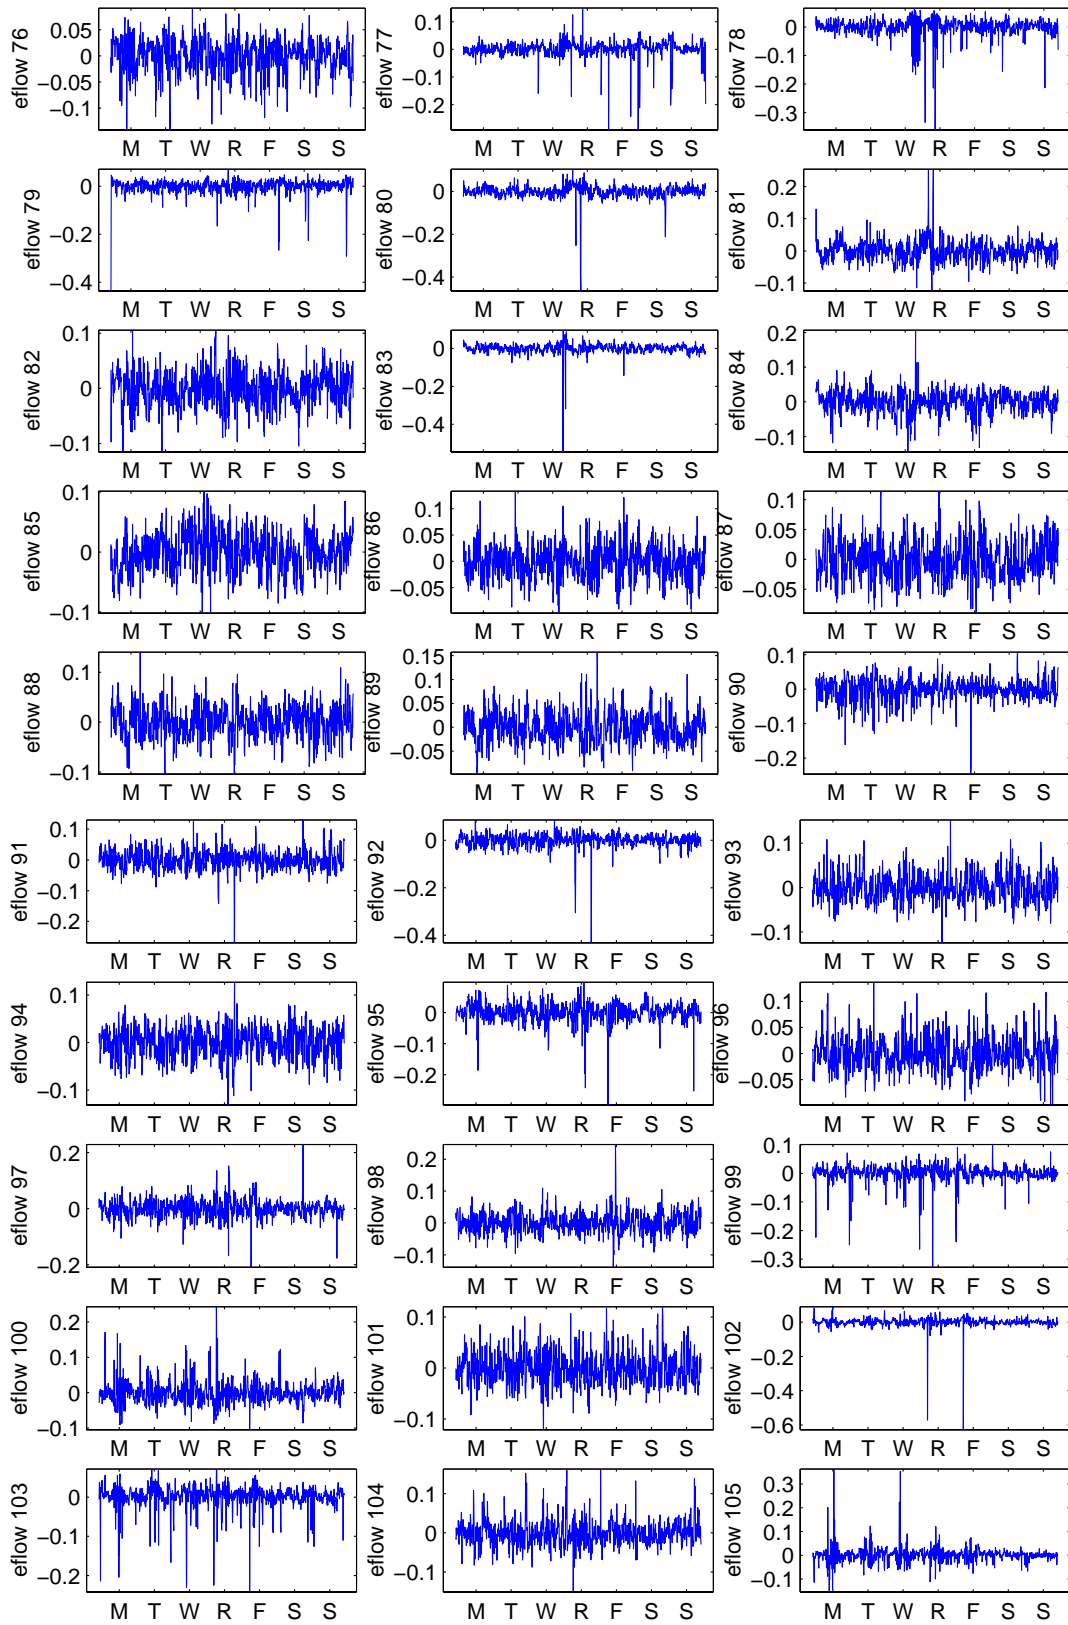
In this section, we first present all the 169 eigenflows from the Sprint-1 dataset. Following this, we present the entire set of 121 eigenflows from the Abilene dataset.

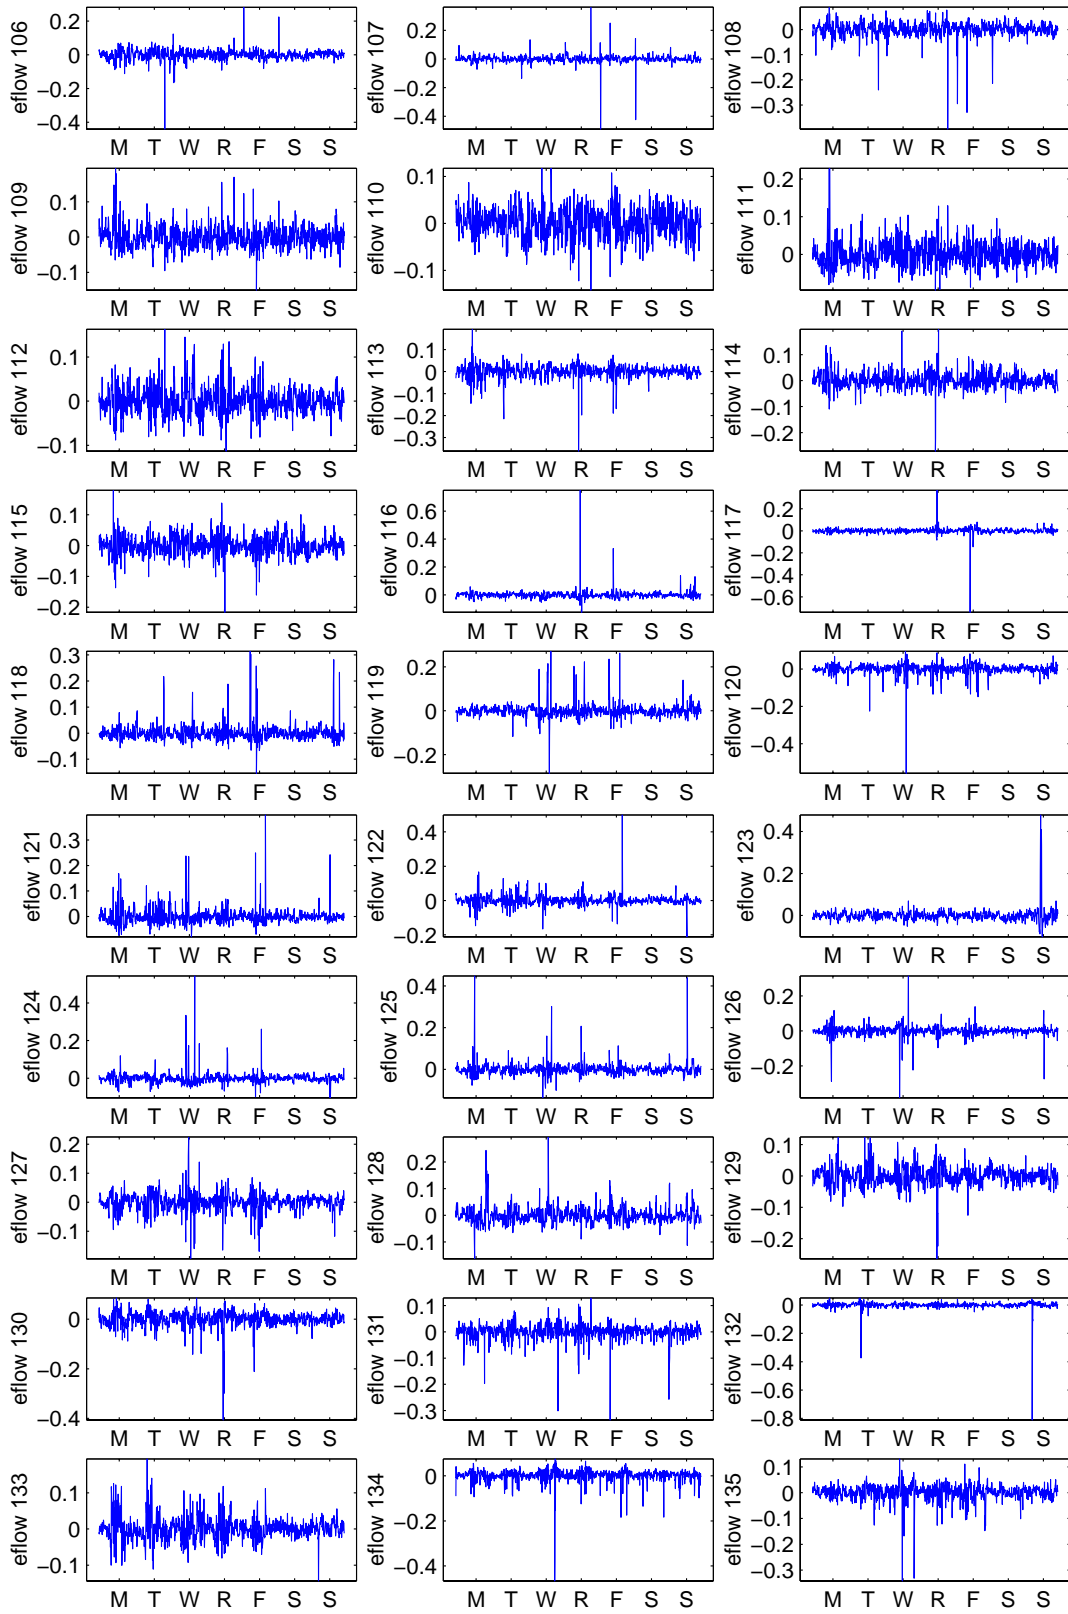
2.1 Sprint-1

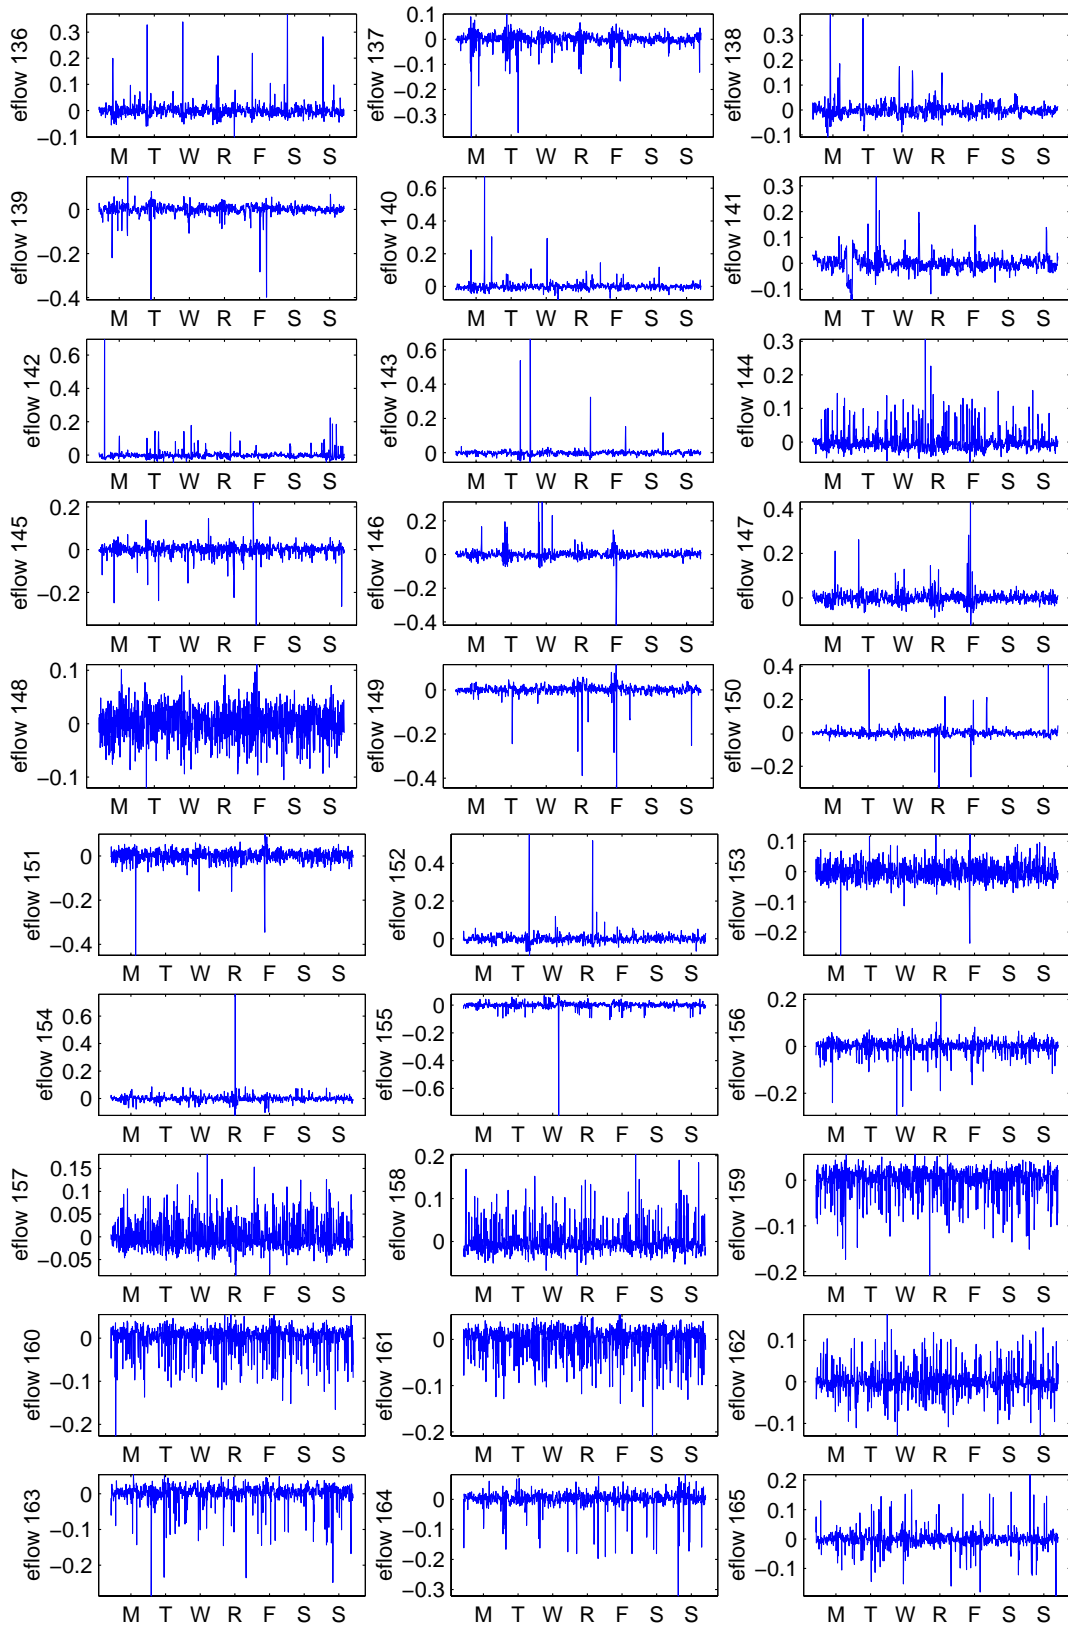







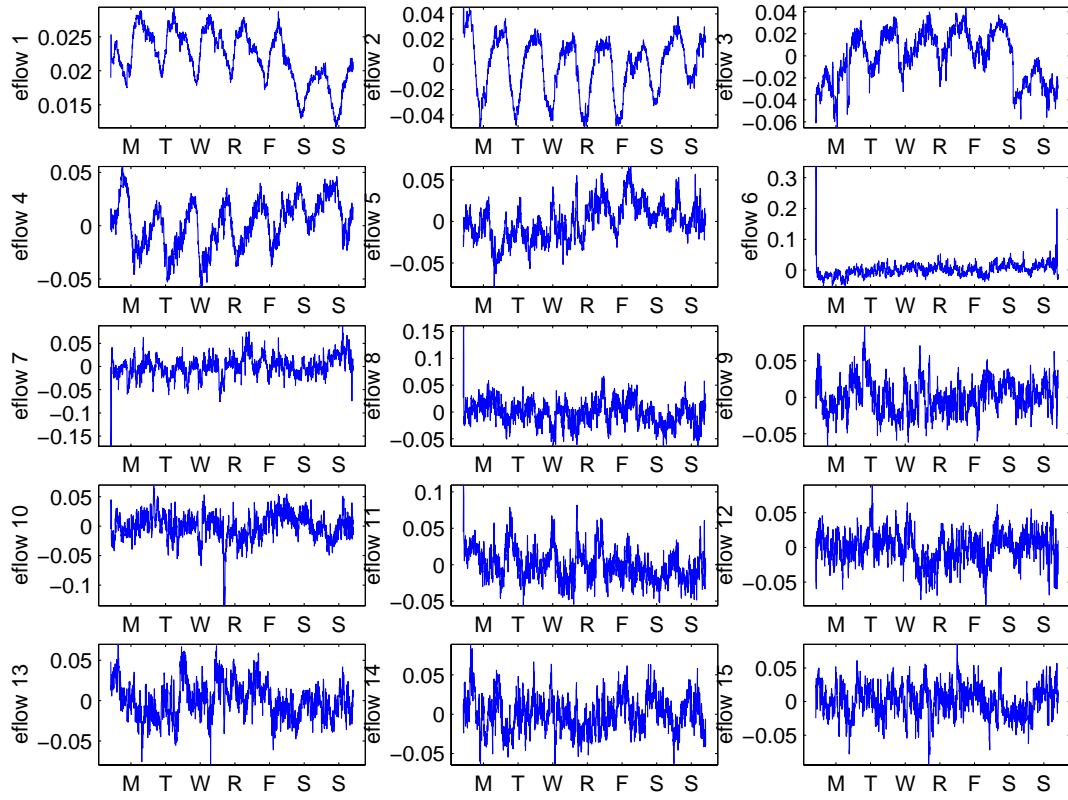


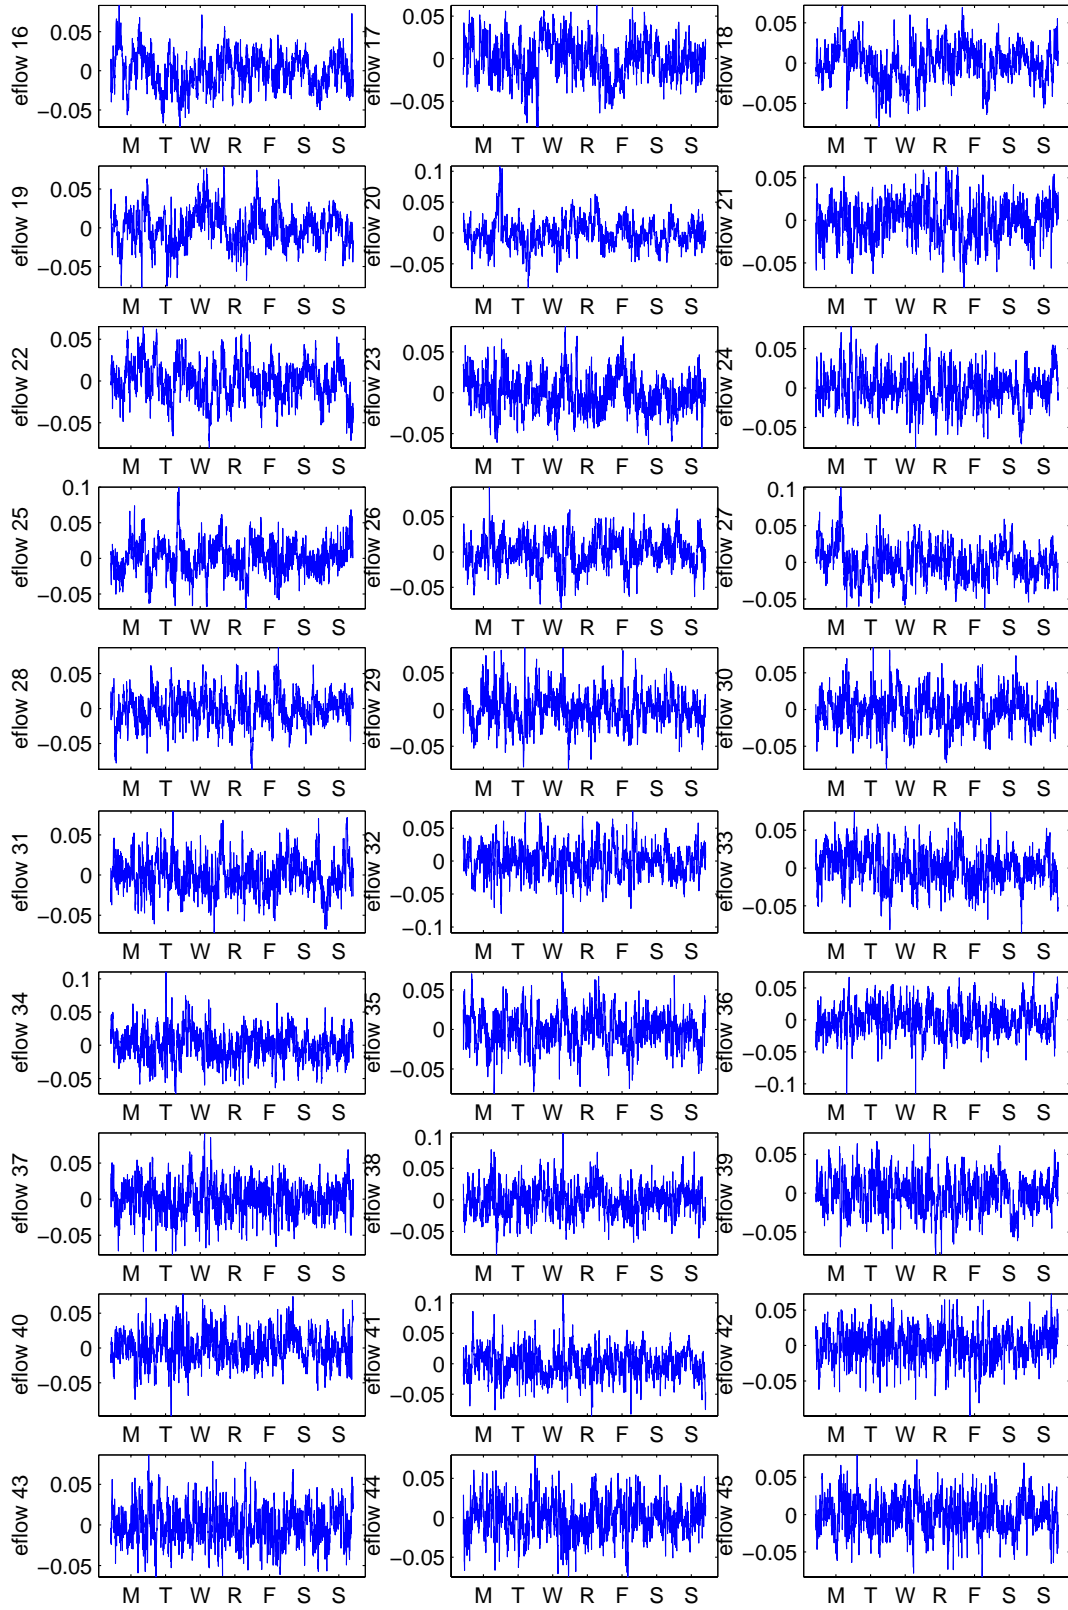


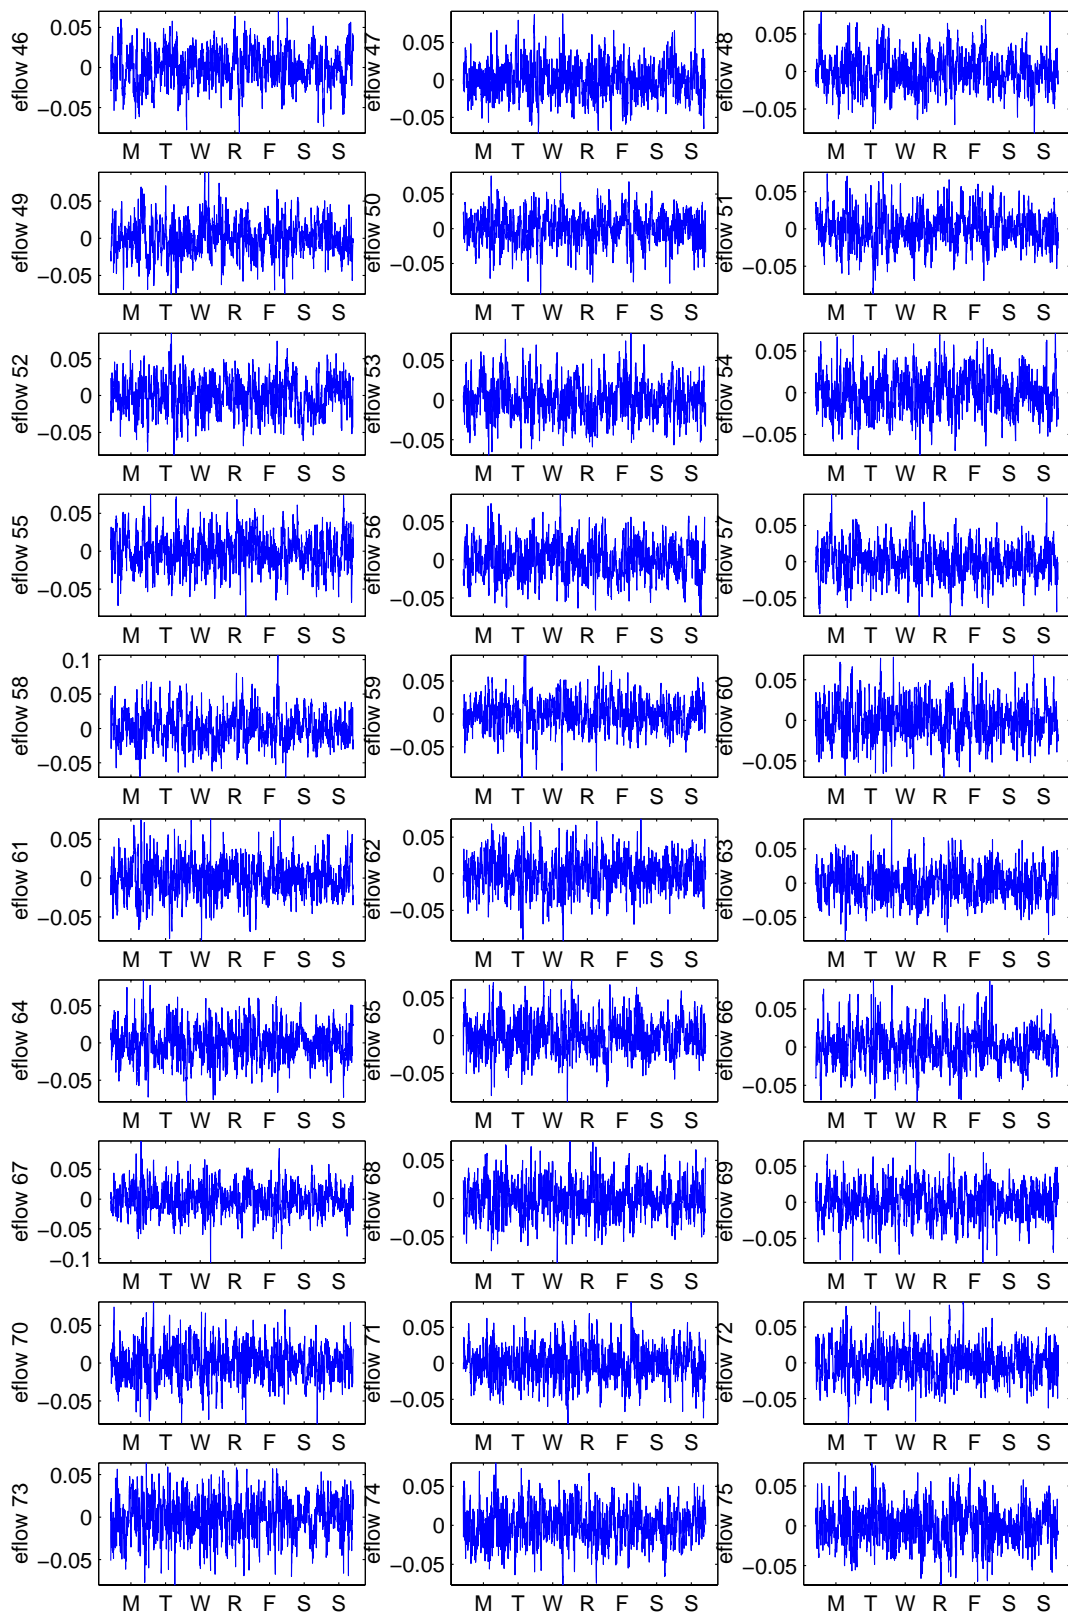


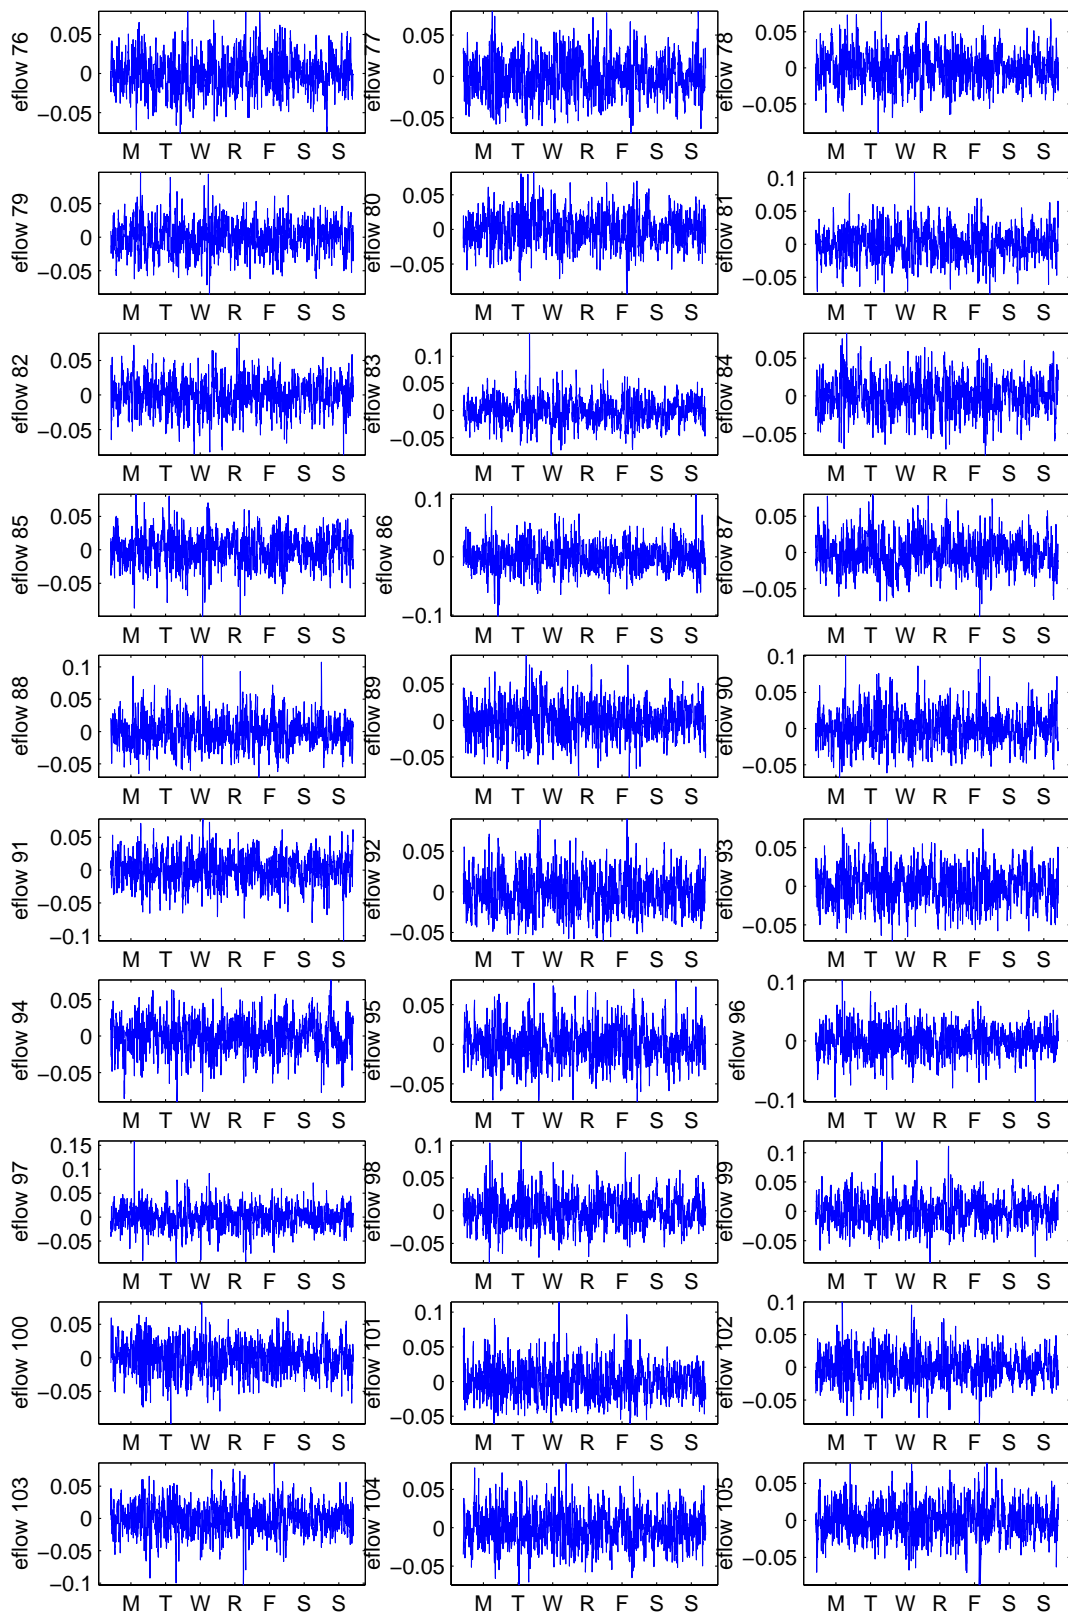


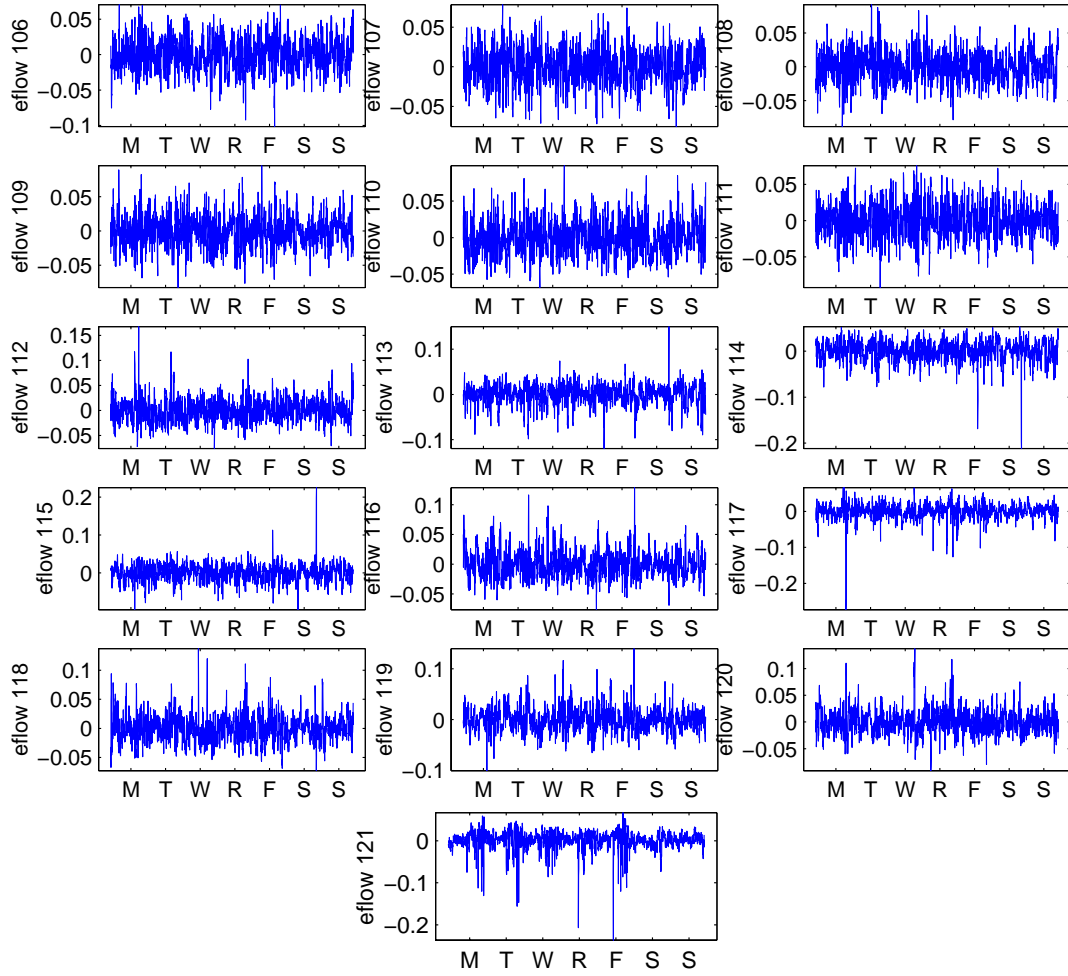
2.2 Abilene







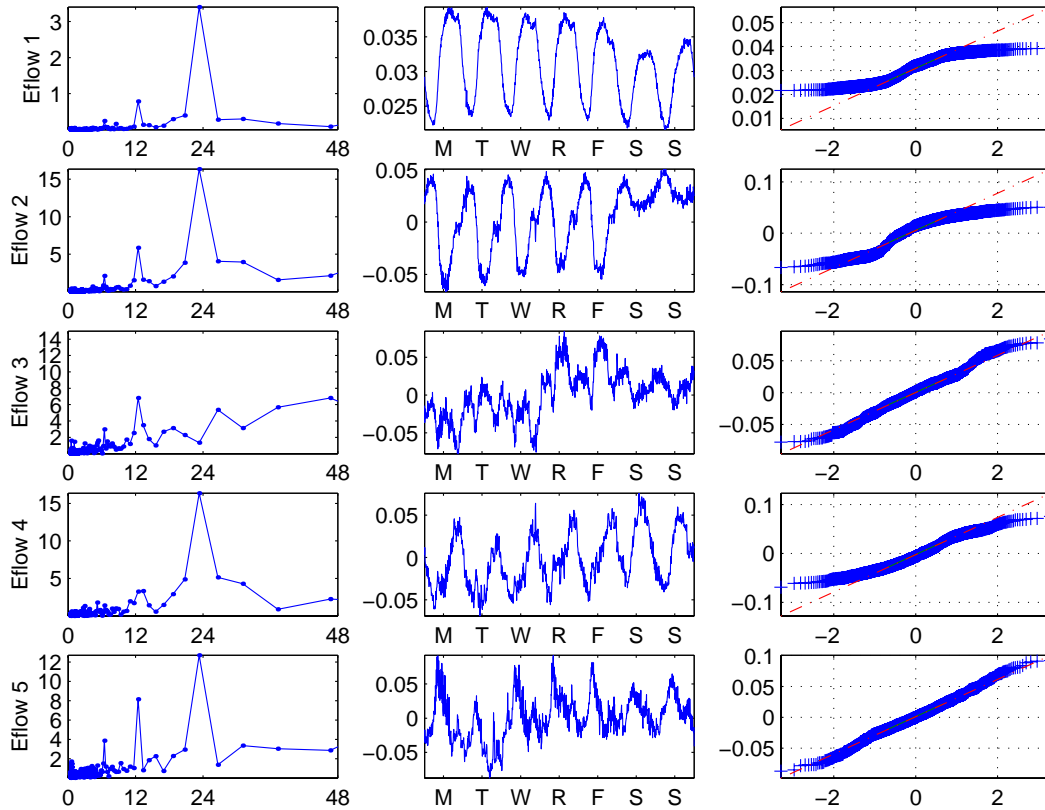


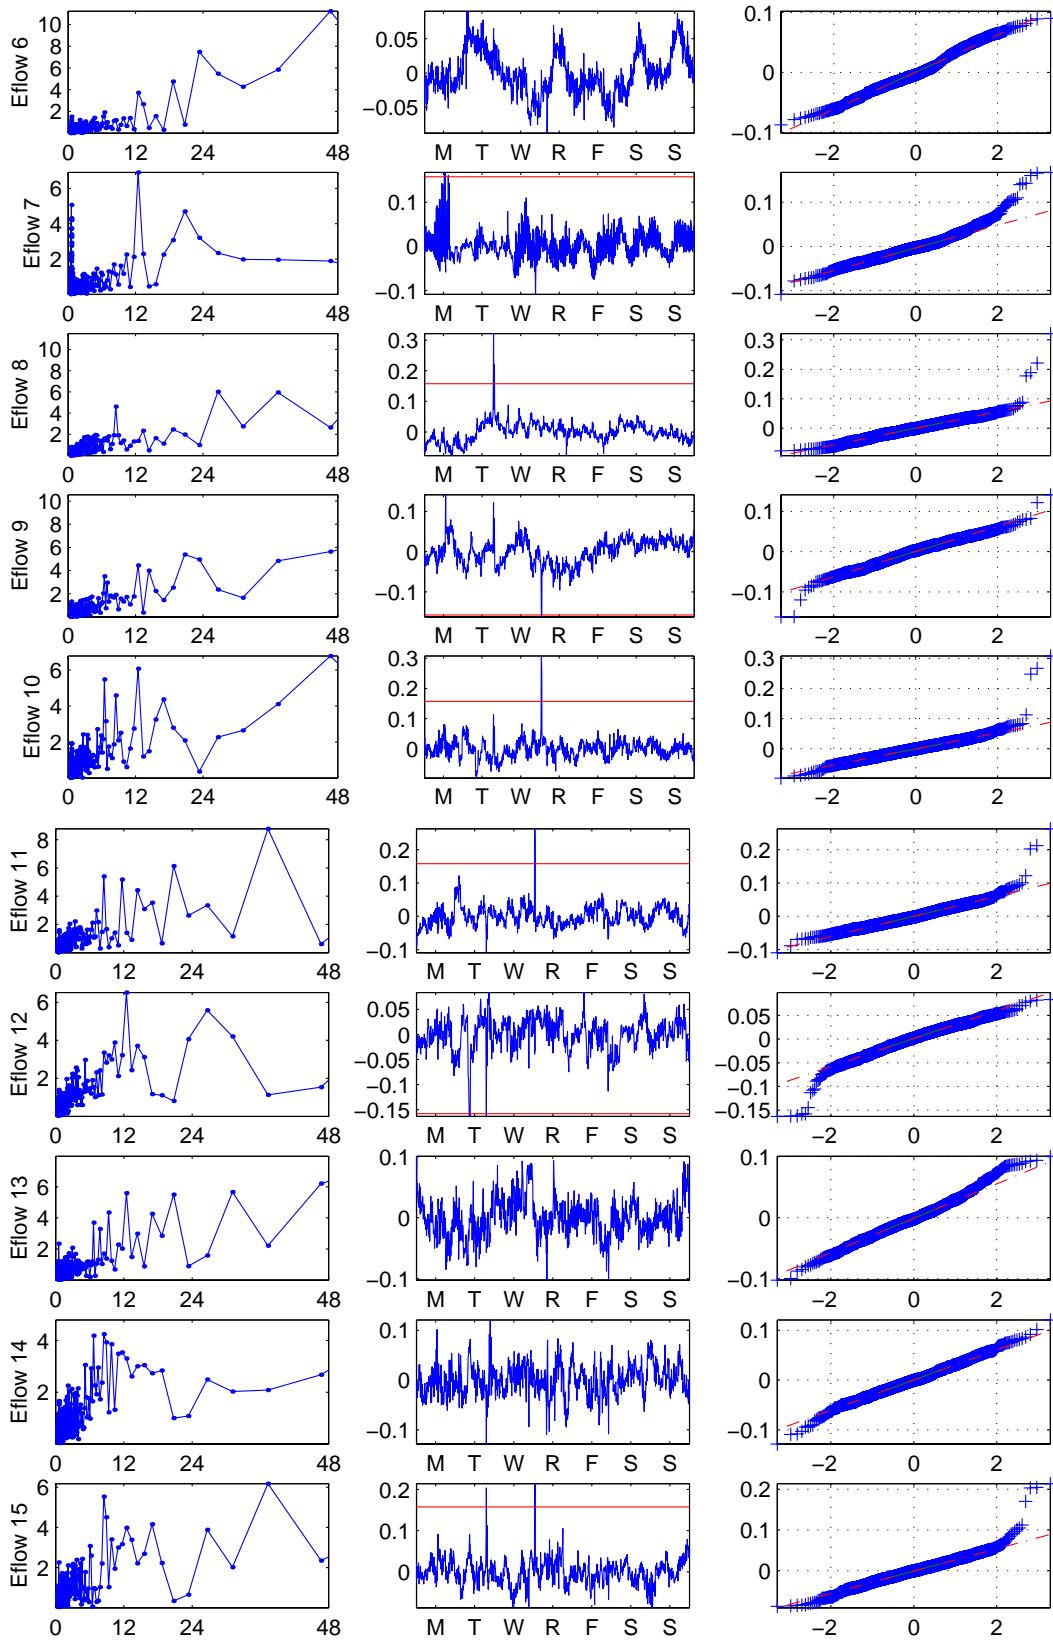


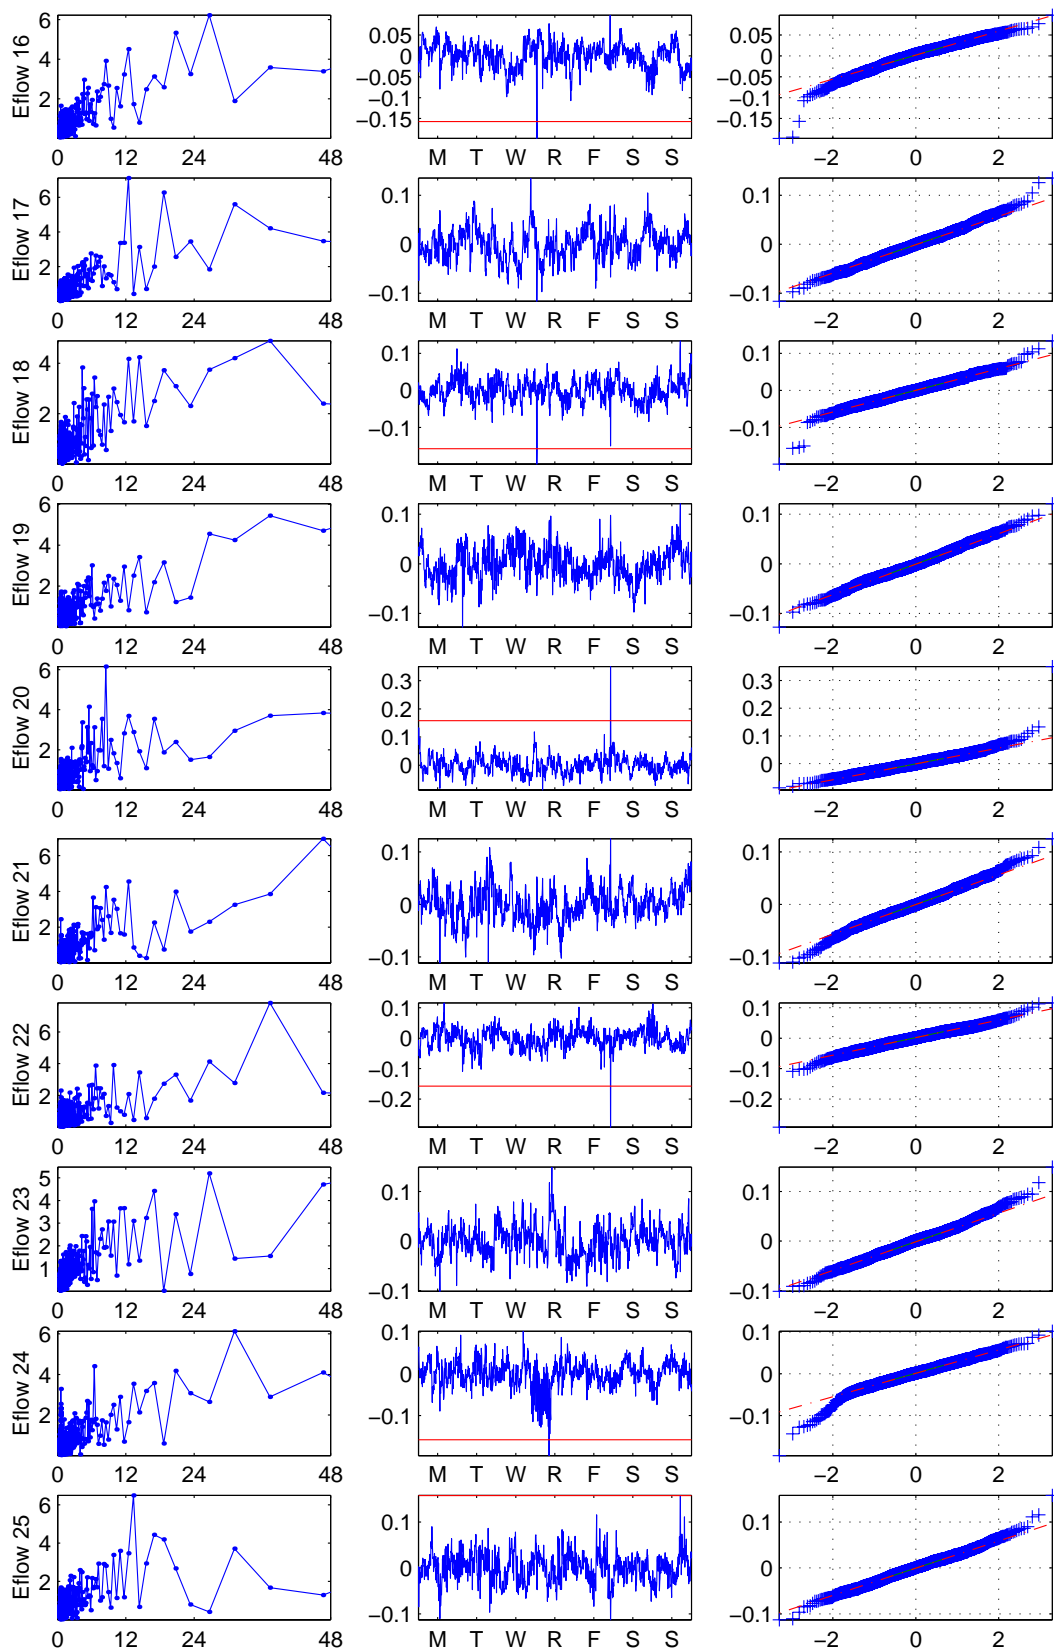
3 Classifying Eigenflows

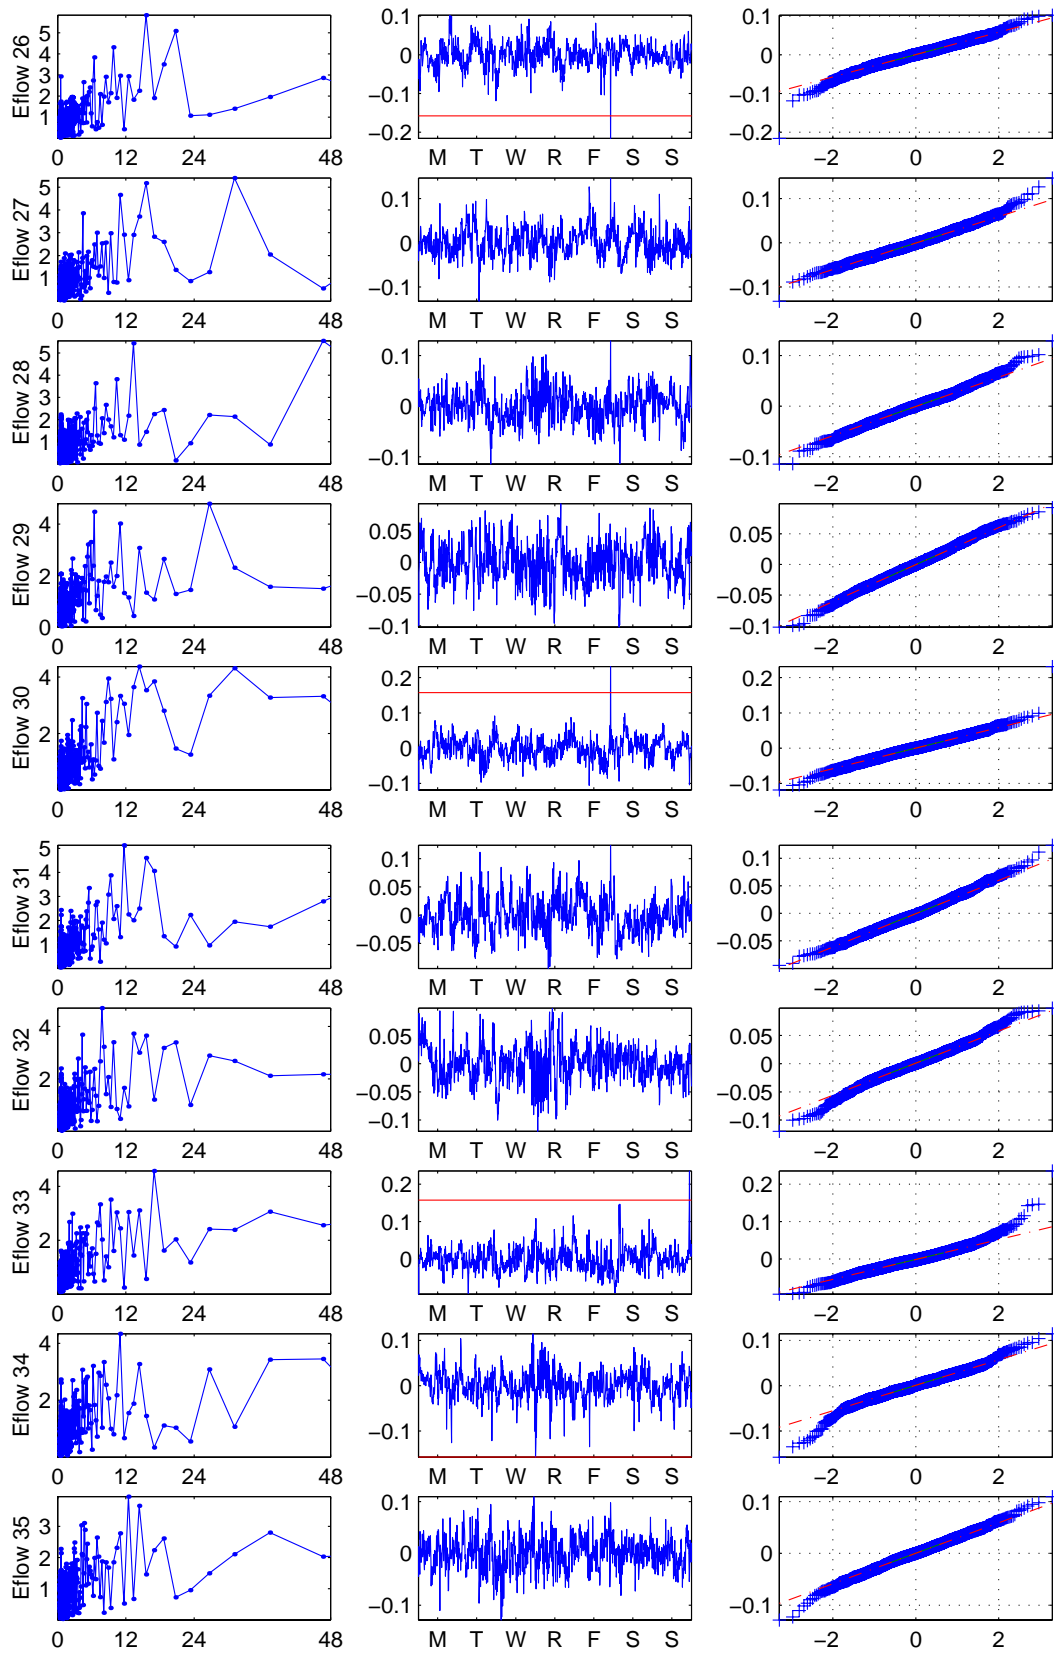
In this section, we present the results of the procedure detailed in our original paper for classifying eigenflows. For all the 169 and 121 eigenflows from the Sprint-1 and Abilene datasets respectively, we plot (a) Fourier Frequency, (b) 5σ threshold and (c) QQplot. Recall that 99.9% of the total energy is captured by the first 50 and the first 90 eigenflows for Sprint-1 and Abilene respectively. However, here we present the entire set of eigenflows from both datasets, for completeness. Note that the 5σ threshold lines (in red) are only plotted if the threshold test is violated.

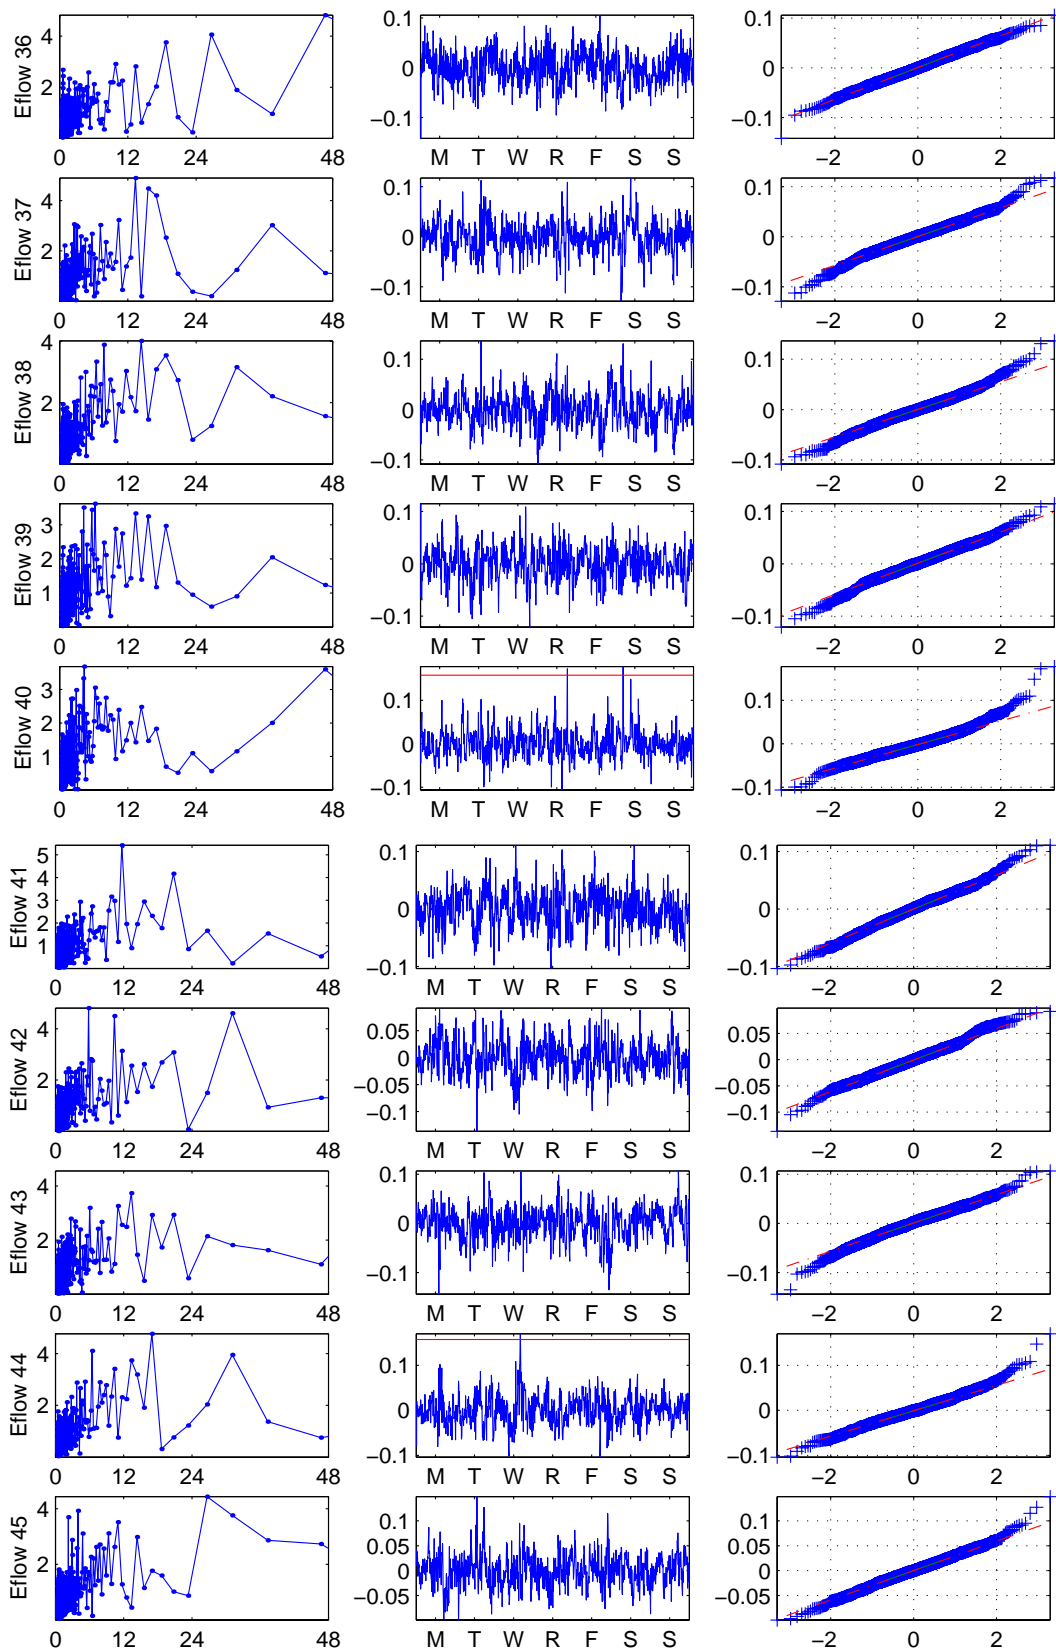
3.1 Sprint-1

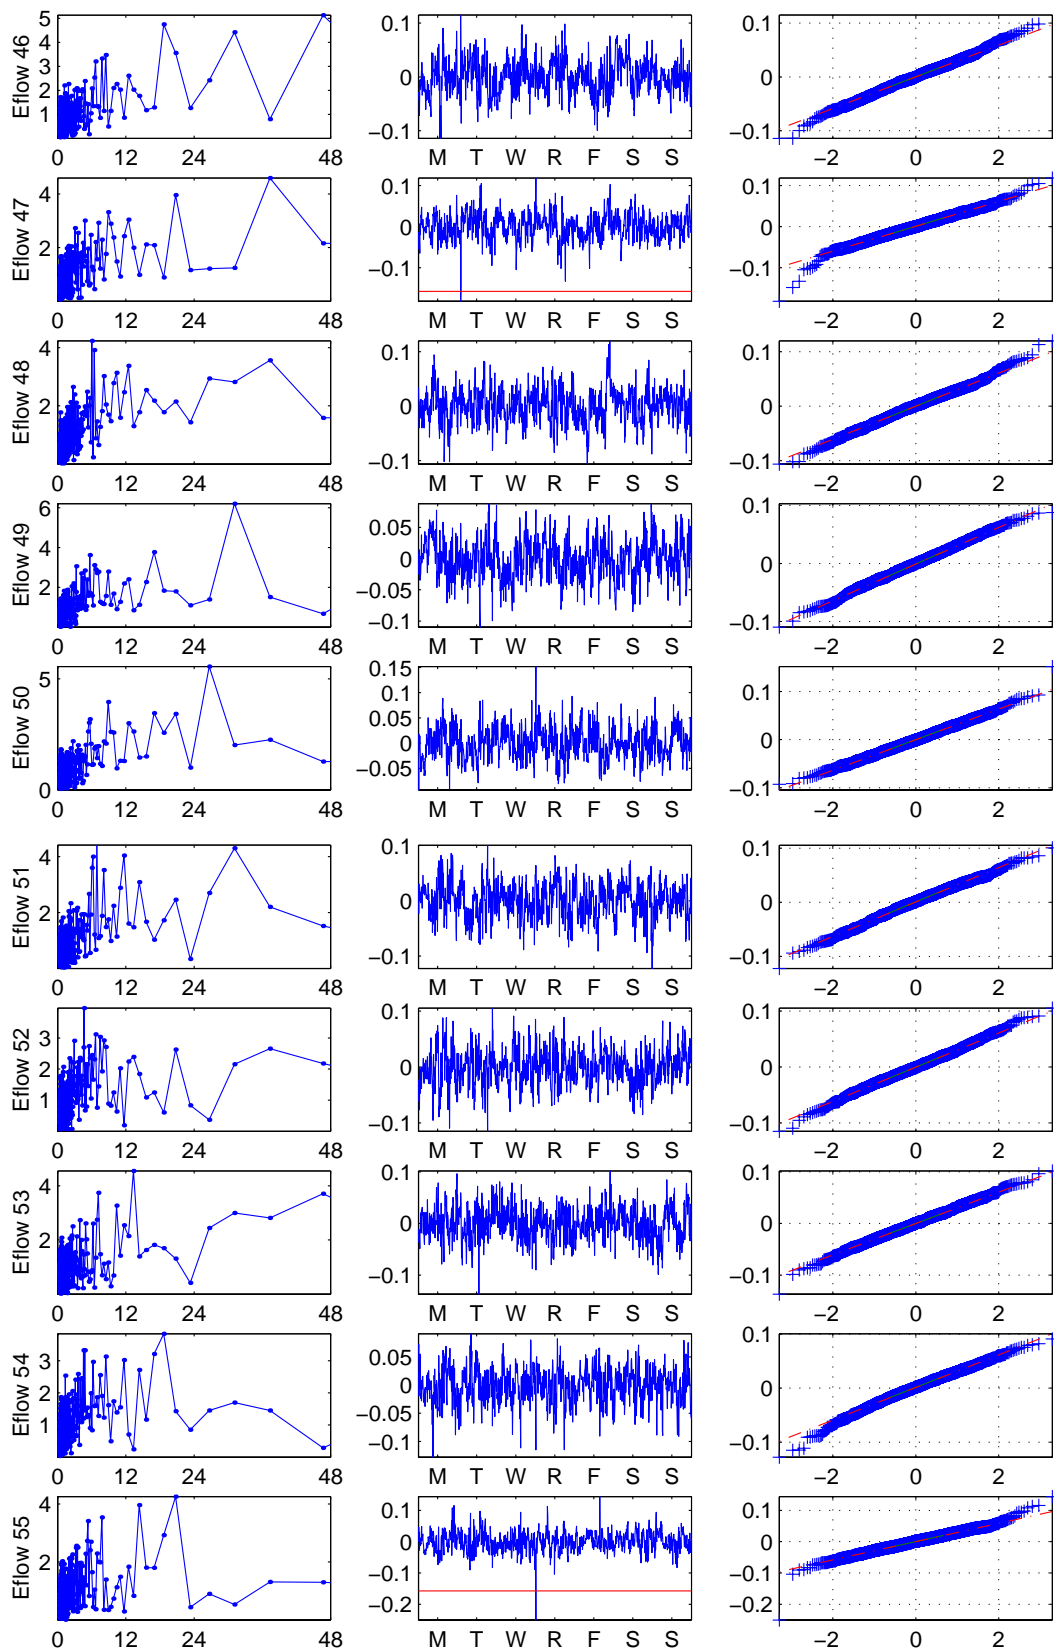


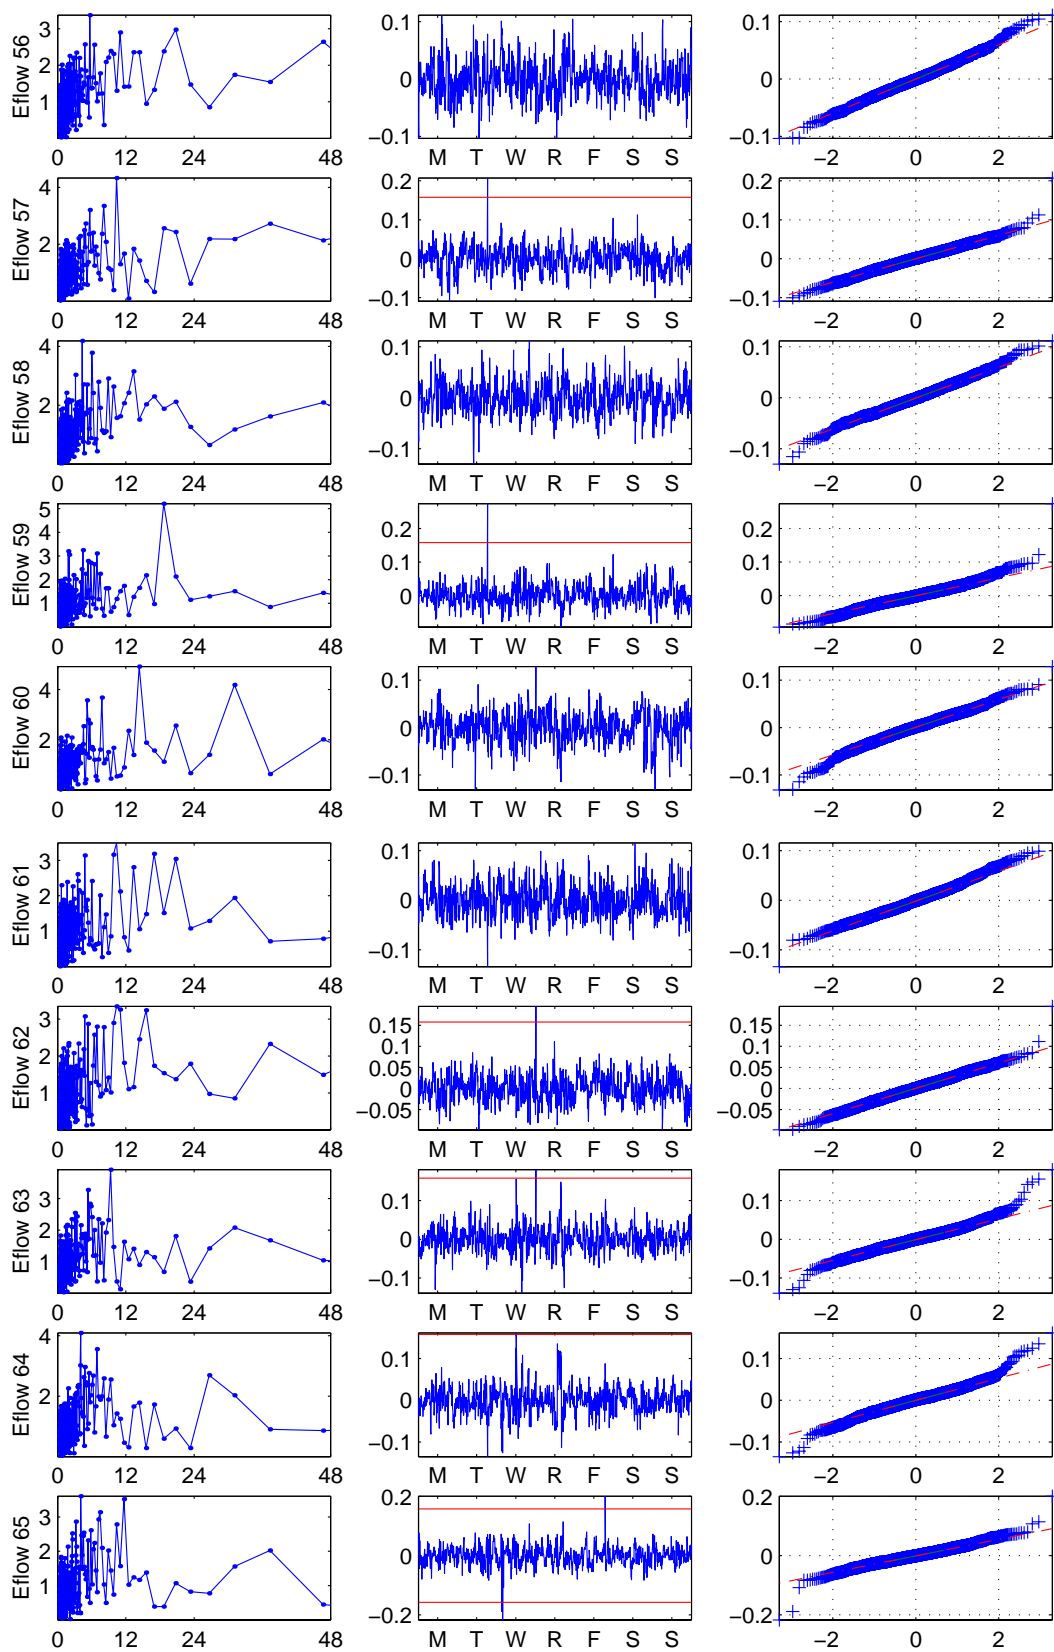


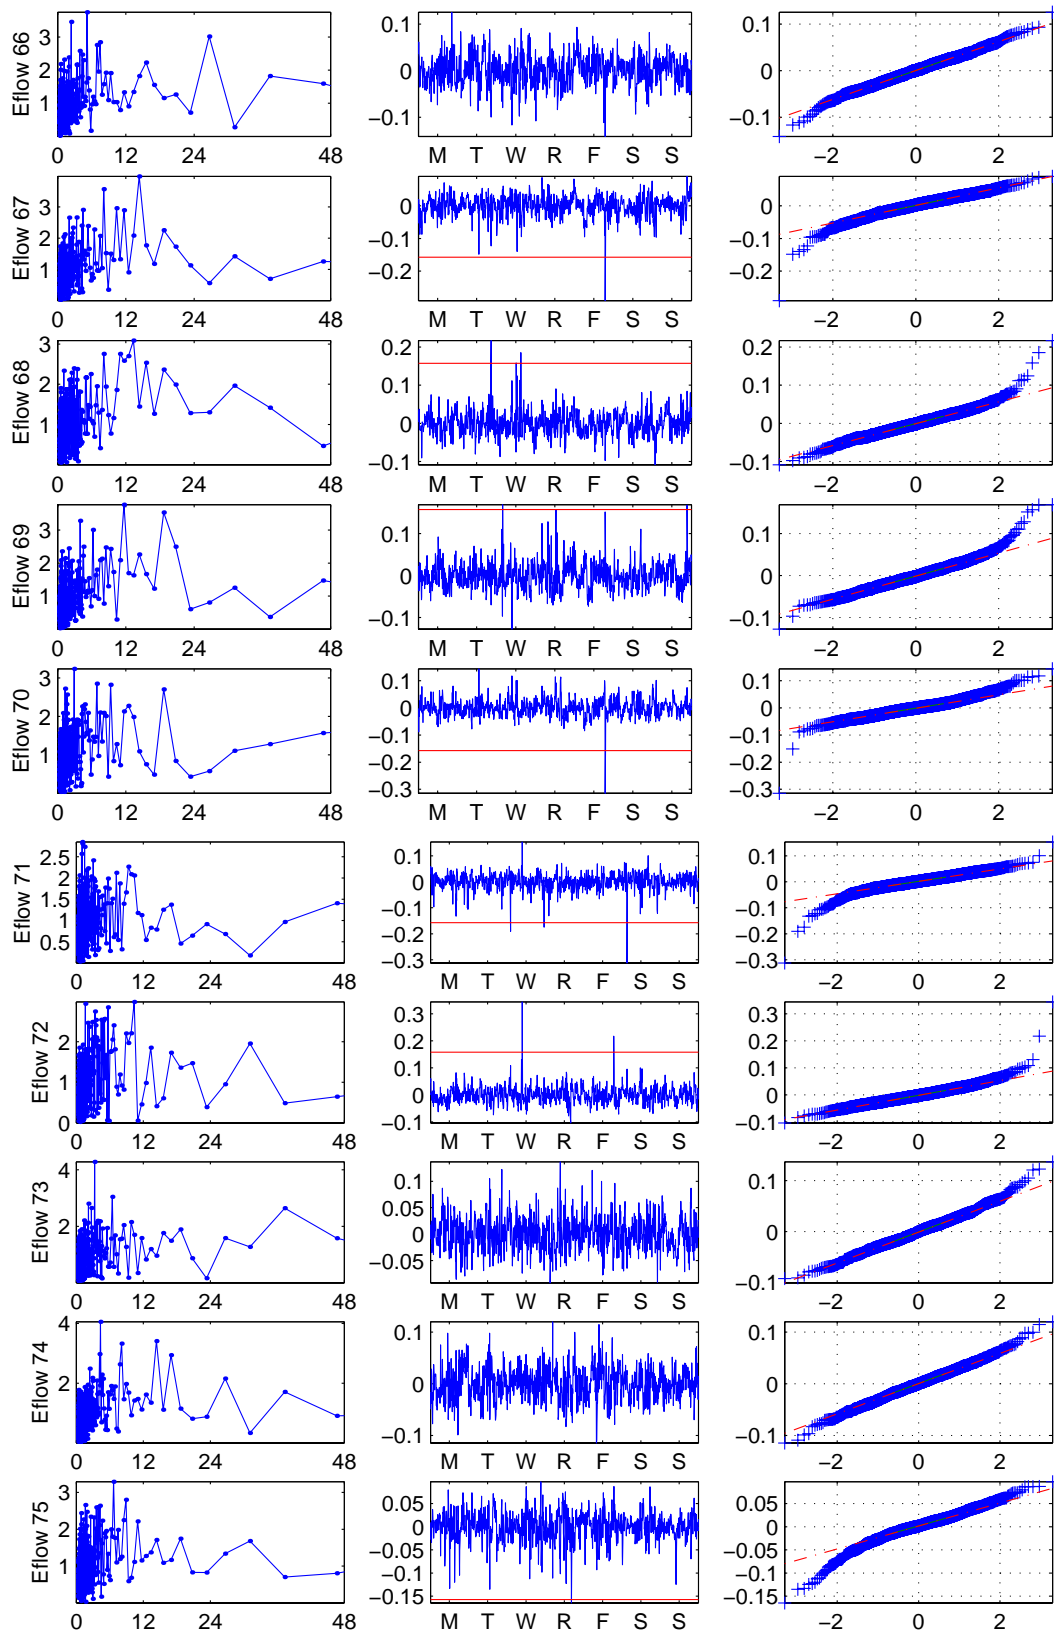


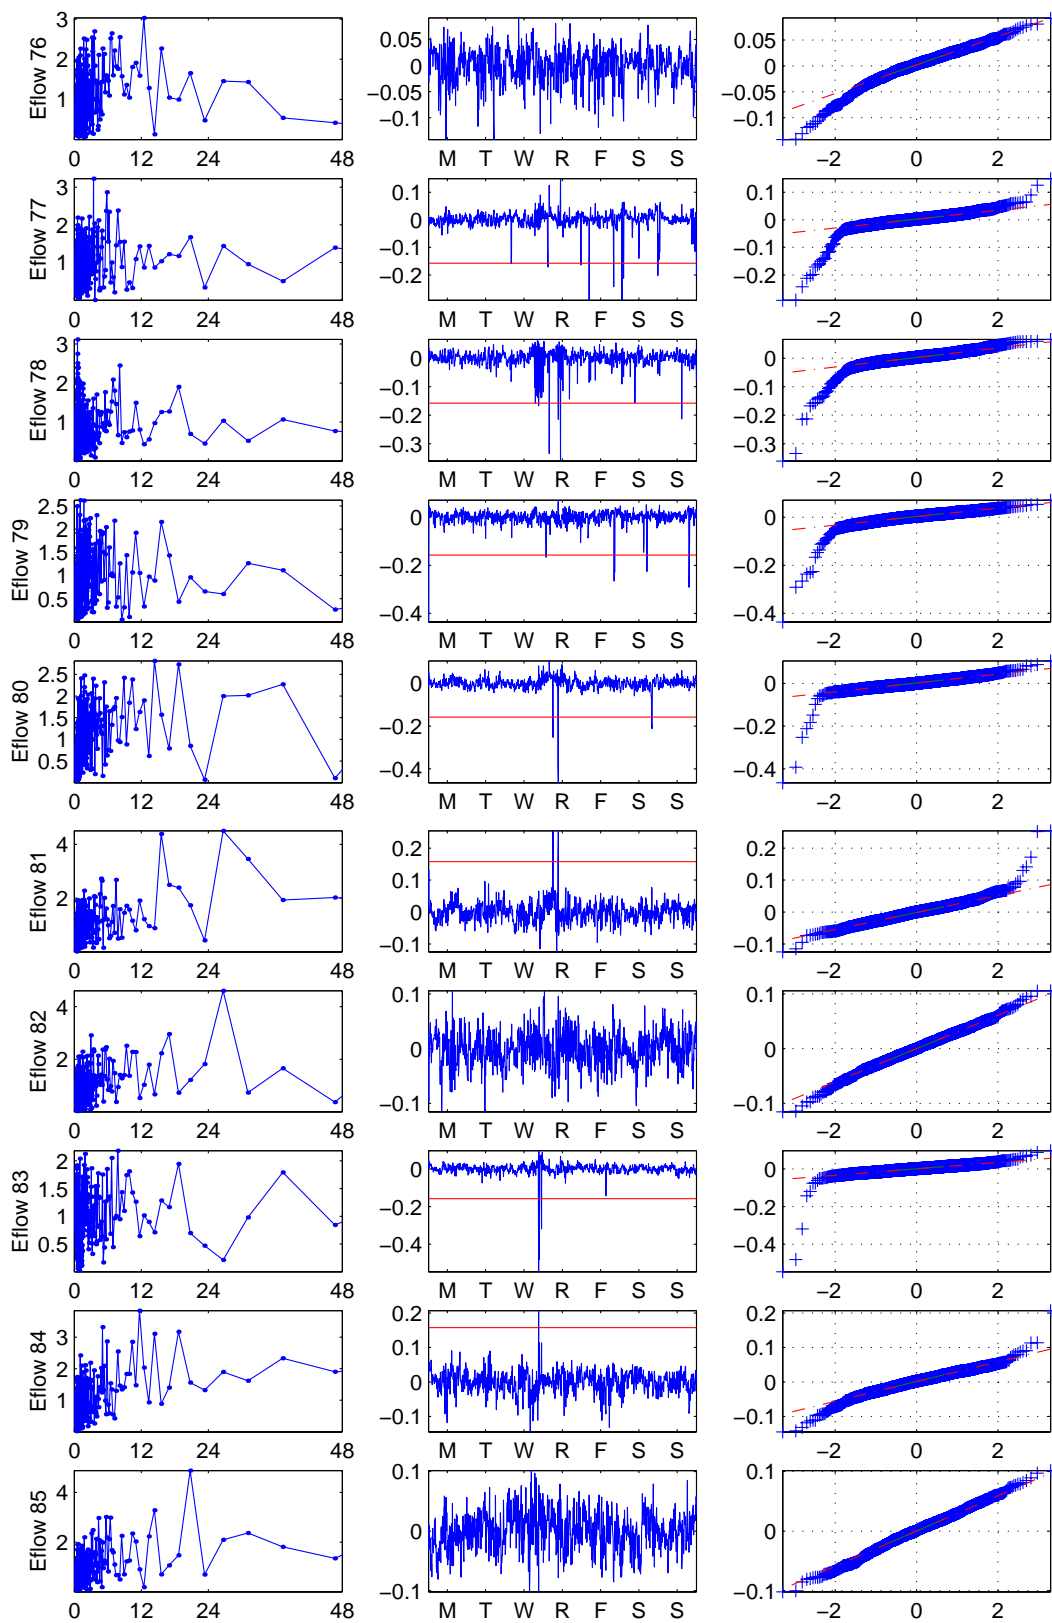


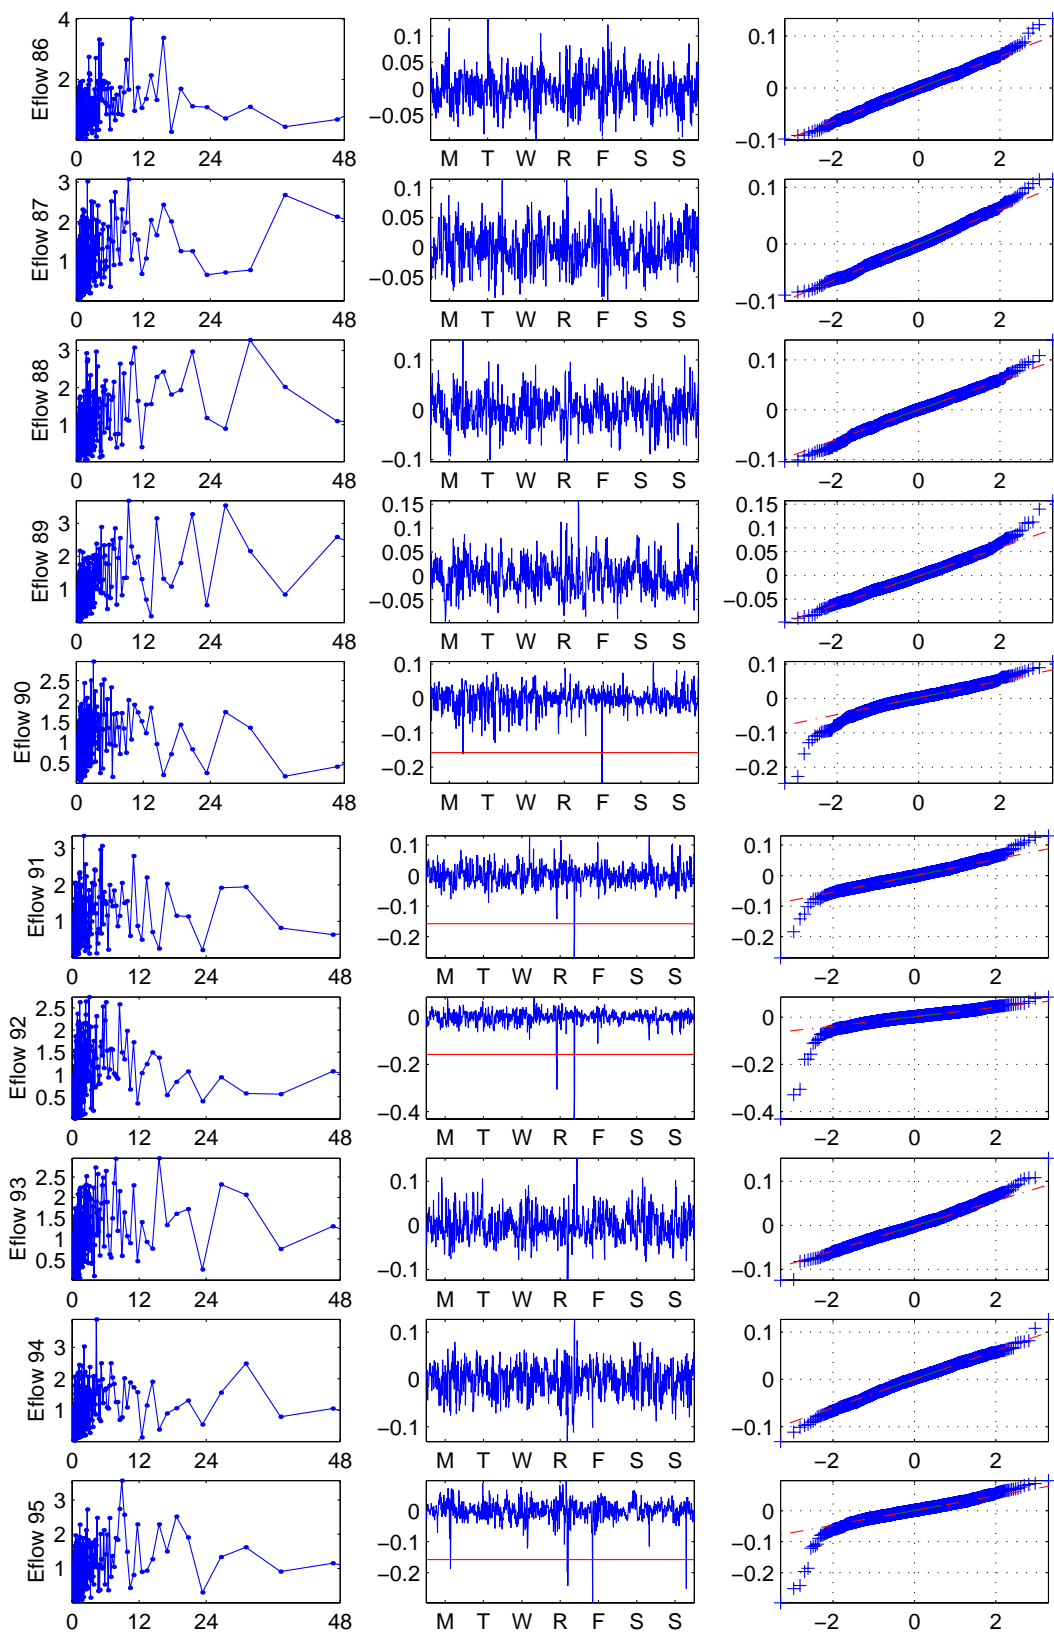


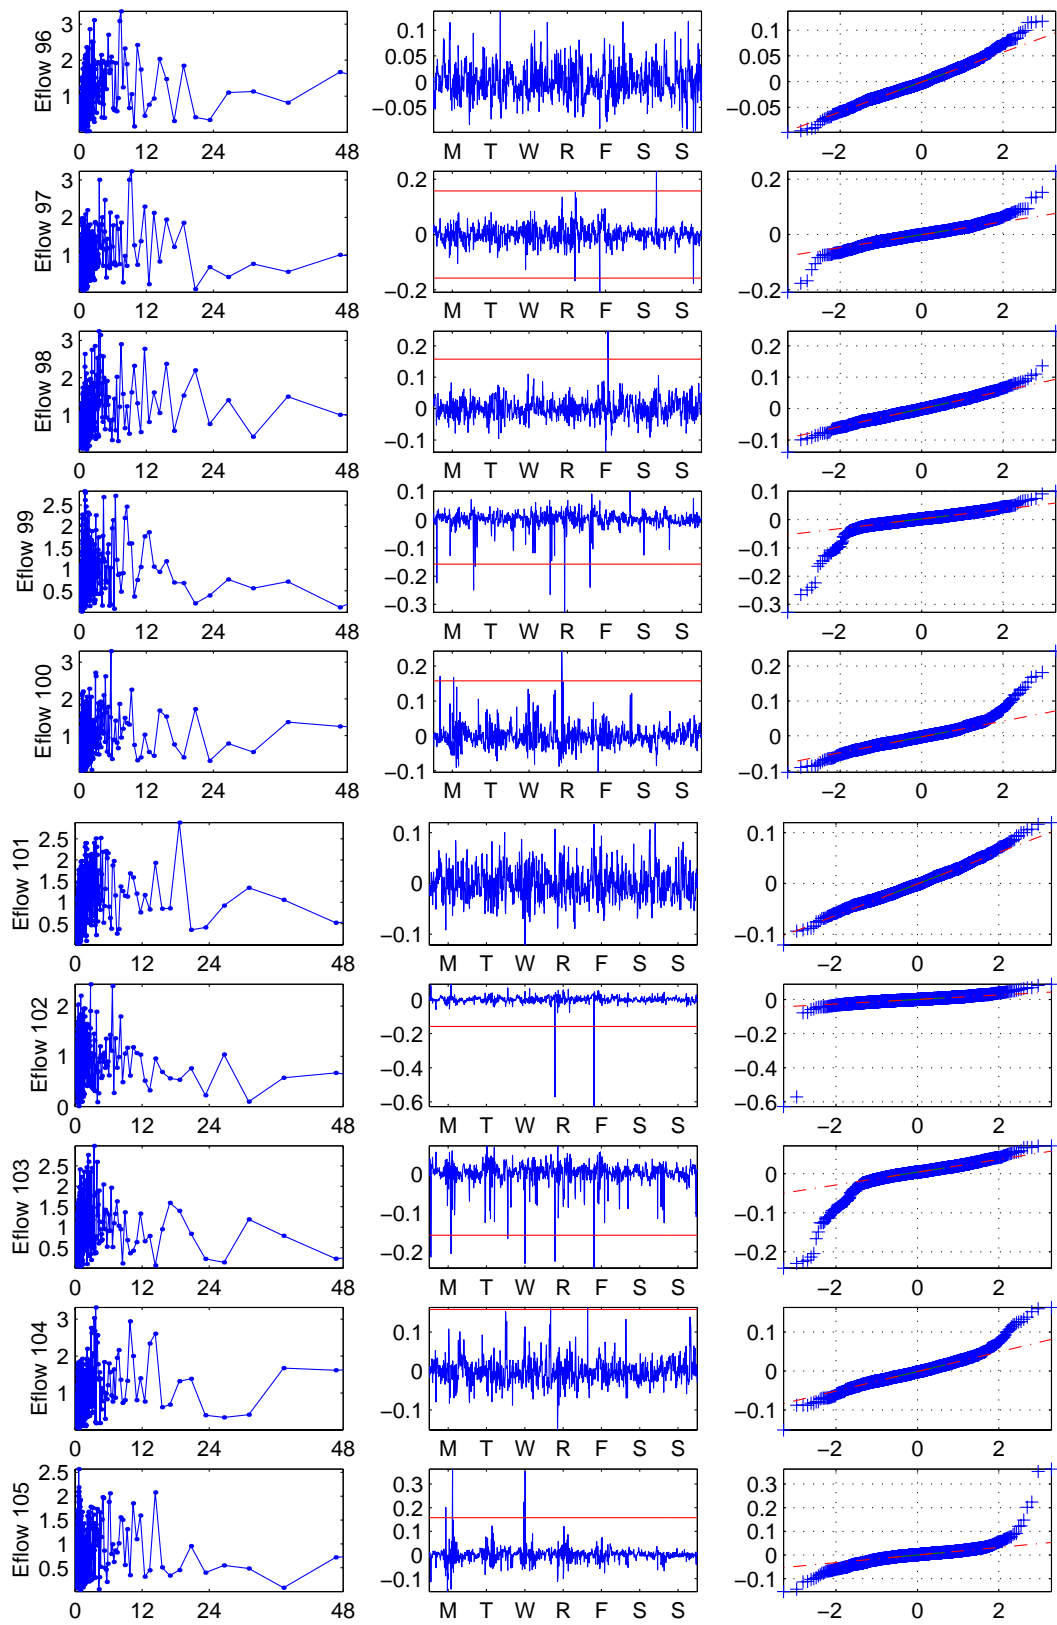


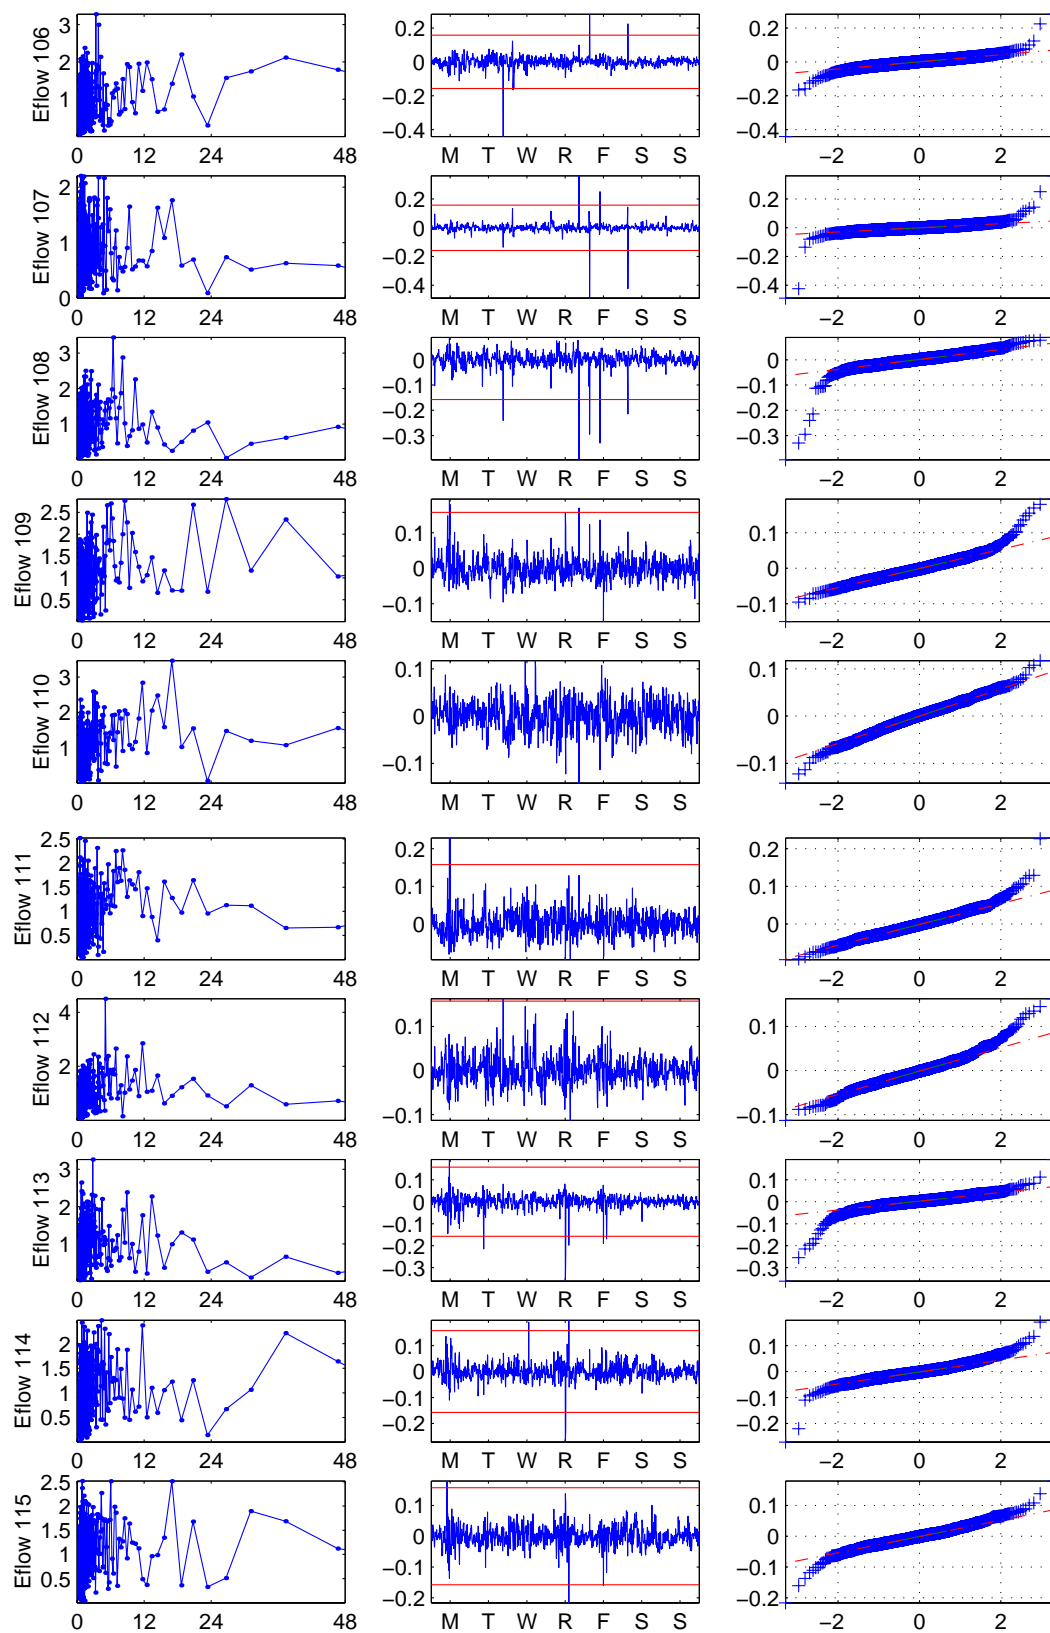


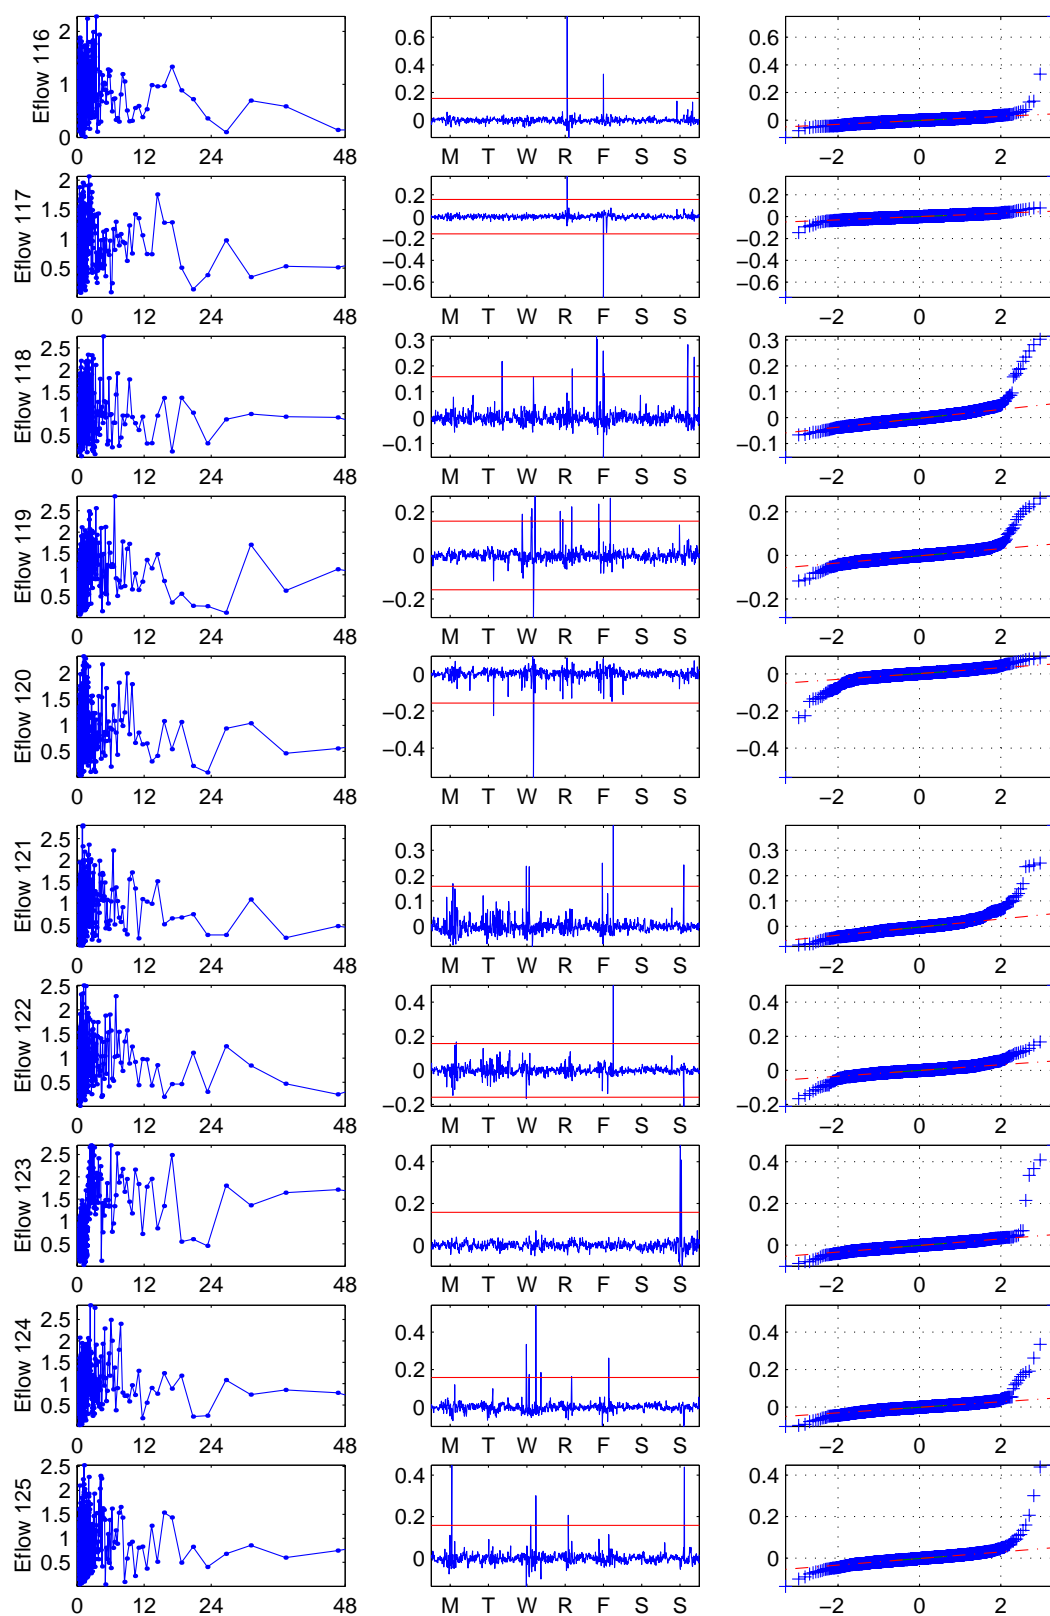


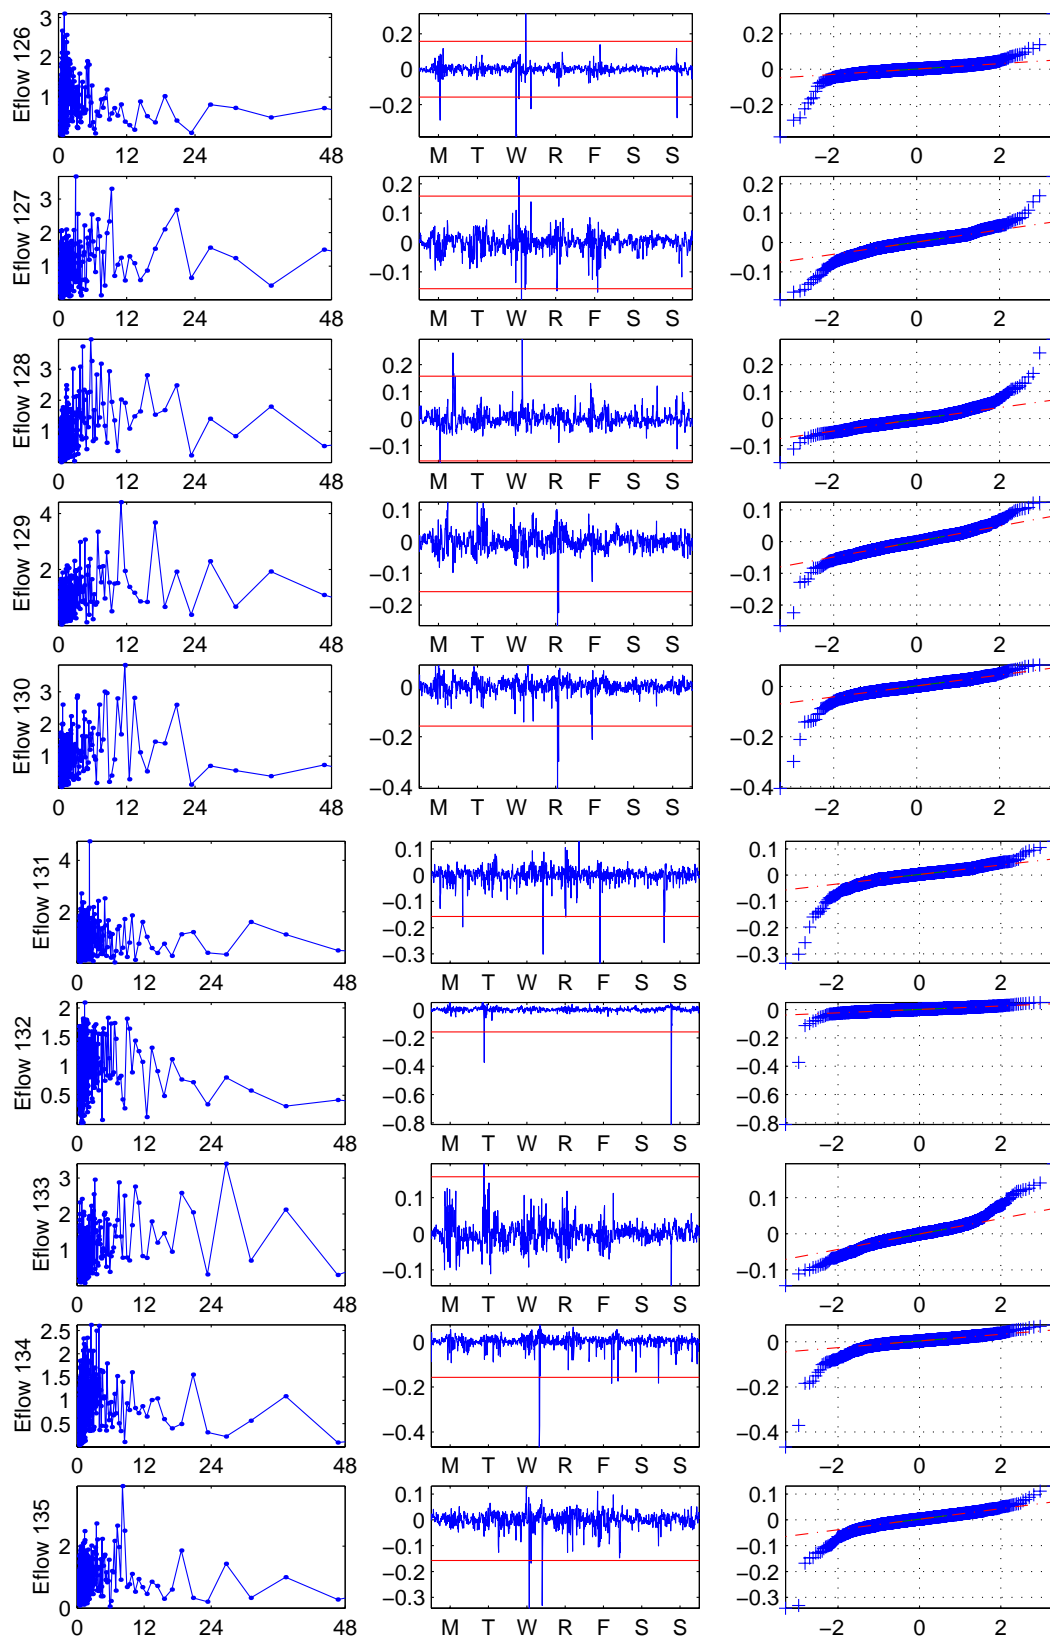


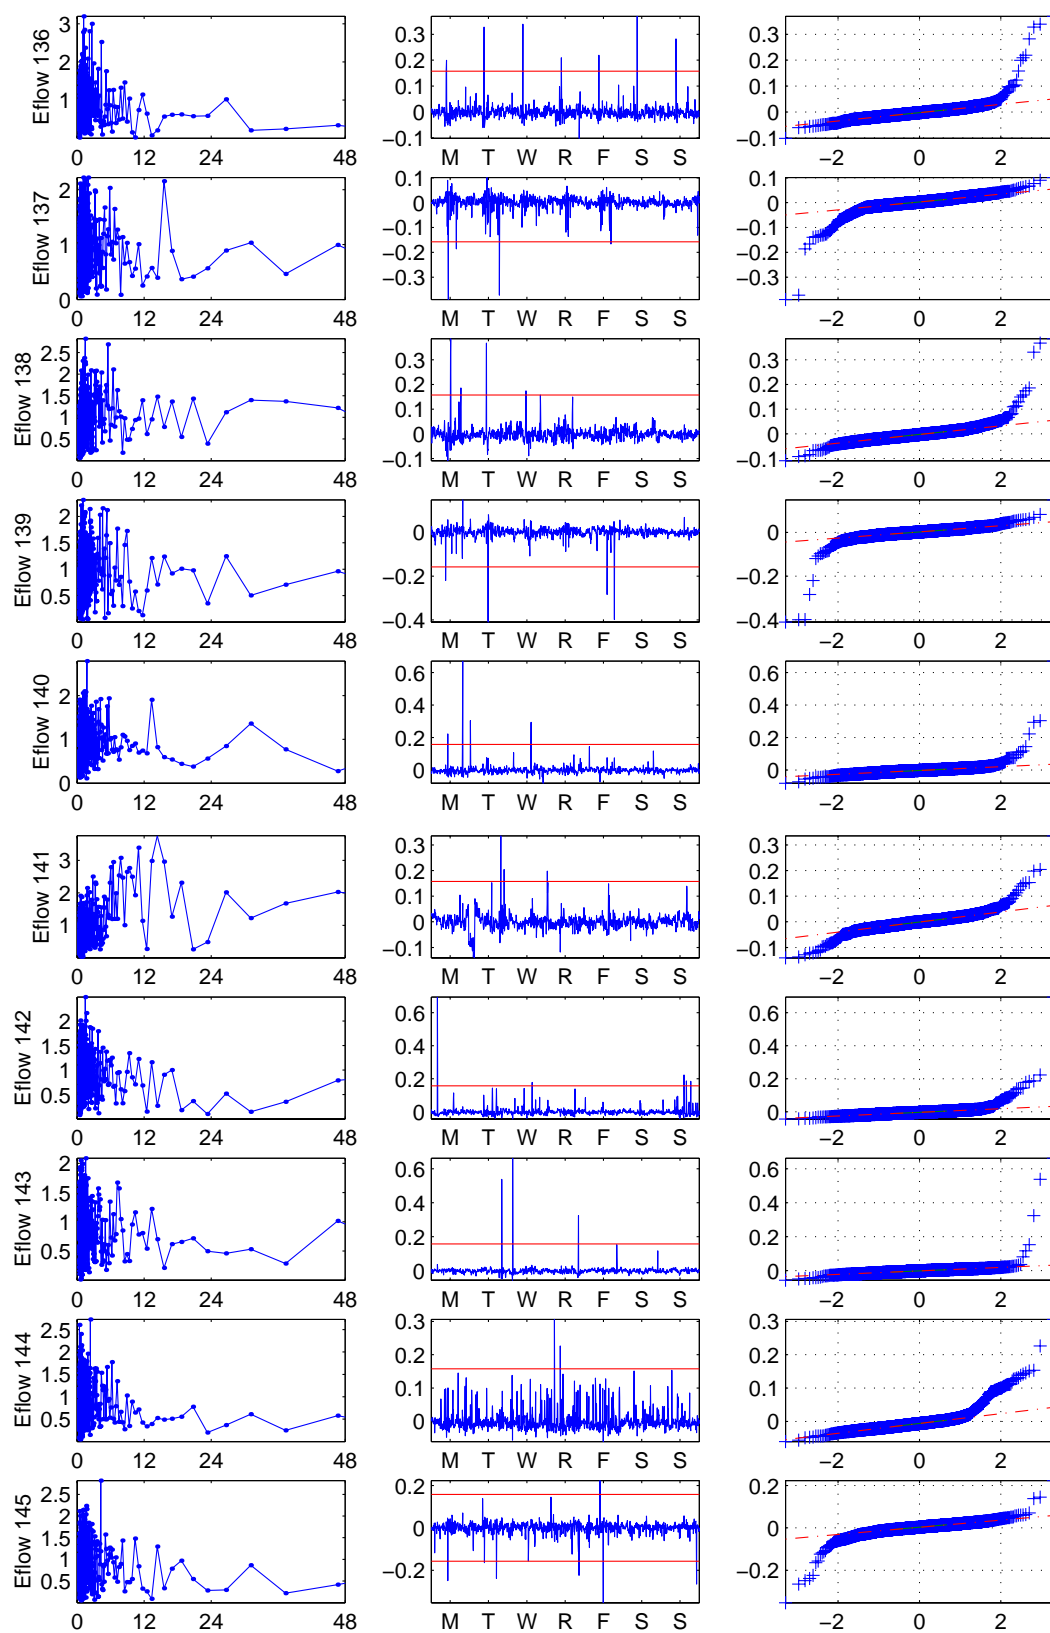


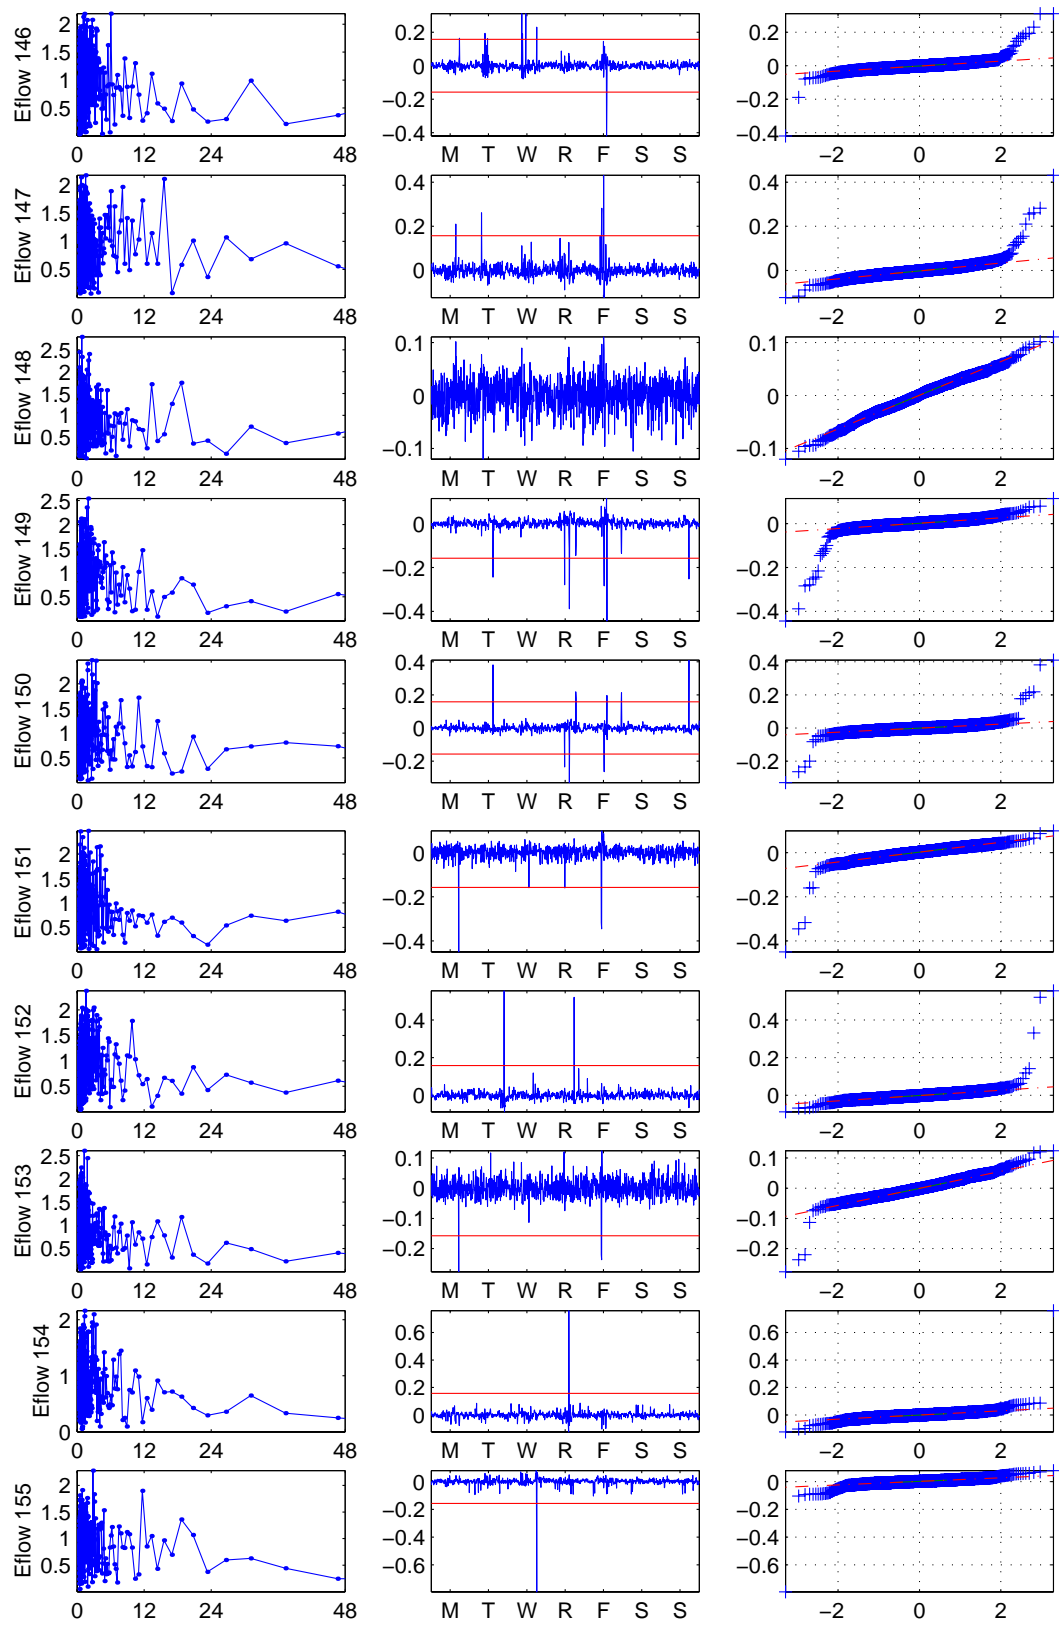


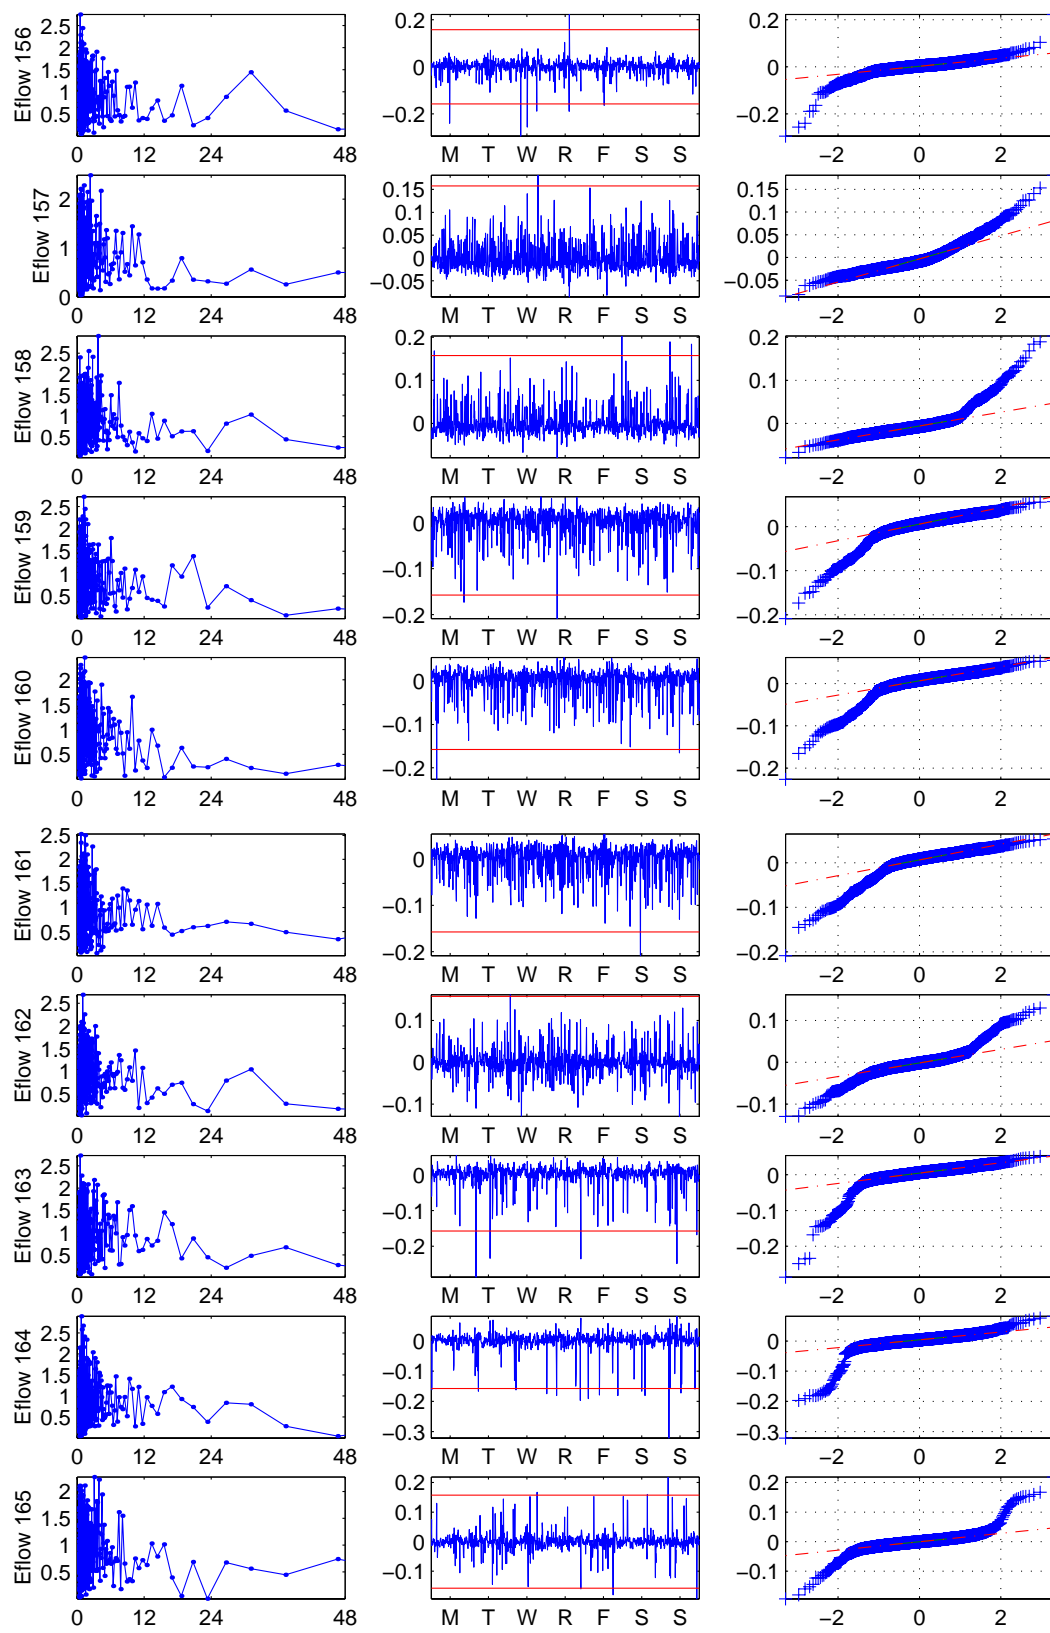


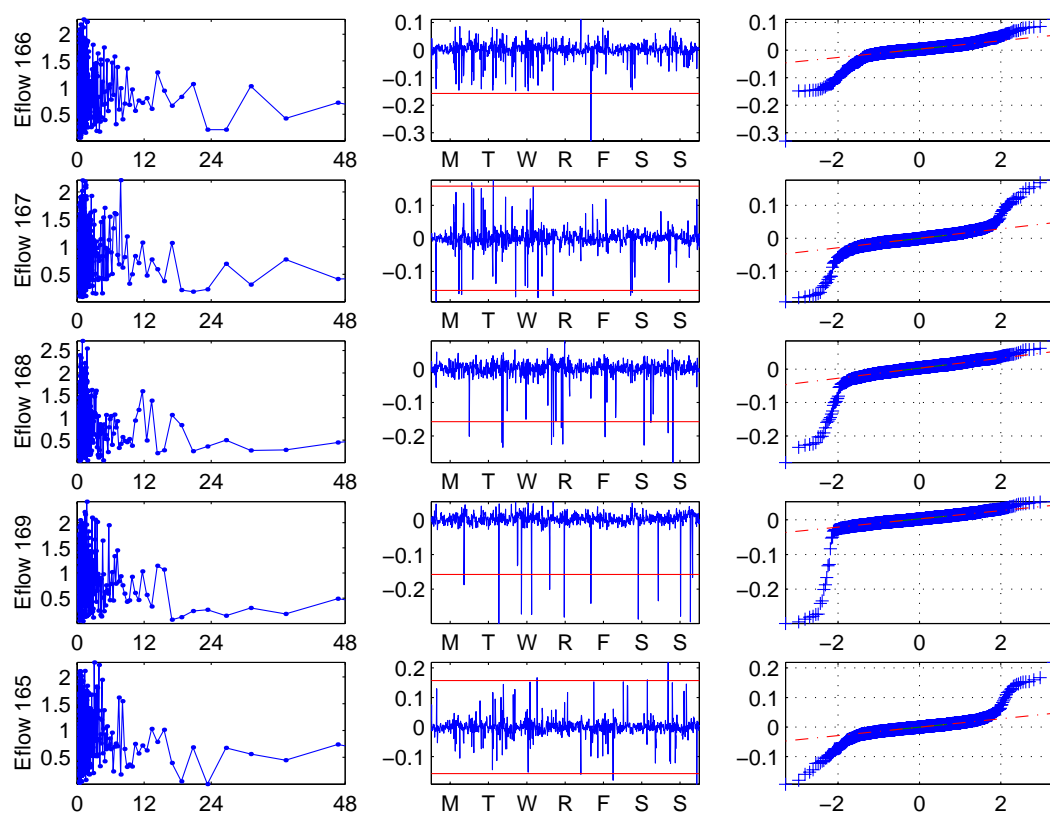












3.2 Abilene

